1. Home . . . 5
1.1 FusionReactor User's Guide . . . 6
1.1.1 Getting Started . . . 9
1.1.2 Requests . . . 10
1.1.2.1 Running Requests . . . 10
1.1.2.2 Request History . . . 14
1.1.2.3 Request Settings . . . 15
1.1.2.4 Request Log (request.log) . . . 17
1.1.2.5 Request Capture . . . 20
1.1.3 Enterprise . . . 20
1.1.3.1 Enterprise Dashboard . . . 20
1.1.3.2 Manage Servers . . . 26
1.1.3.3 Manage Groups . . . 28
1.1.3.4 Enterprise Settings . . . 28
1.1.3.5 Using Monitored Server Scripts . . . 32
1.1.4 Metrics . . . 37
1.1.4.1 System Metrics . . . 37
1.1.4.2 Longest Requests . . . 38
1.1.4.3 Slow Requests . . . 40
1.1.4.4 Longest JDBC . . . 43
1.1.4.5 Metrics Settings . . . 44
1.1.5 Resources . . . 45
1.1.5.1 Request Activity Graph . . . 46
1.1.5.2 Request Time Graph . . . 47
1.1.5.3 JDBC Activity Graph . . . 48
1.1.5.4 JDBC Time Graph . . . 49
1.1.5.5 Memory Graphs . . . 50
1.1.5.5.1 Memory Heap . . . 50
1.1.5.5.2 Heap and Non Heap . . . 51
1.1.5.5.3 Types of Memory Graphs . . . 53
1.1.5.6 CPU Graph . . . 55
1.1.5.7 List All Threads . . . 56
1.1.5.8 Resource Settings . . . 57
1.1.5.9 Resource Log (resource.log) . . . 58
1.1.5.10 Memory Logs . . . 61
1.1.5.10.1 Memory Heap and Non Heap Log (memorysummary.log) . . . 61
1.1.5.10.2 Types of Memory Logs . . . 63
1.1.5.11 CPU Thread Log (cputhread.log) . . . 66
1.1.6 Crash Protection . . . 69
1.1.6.1 Crash Protection Restrictions . . . 69
1.1.6.2 Crash Protection Settings . . . 71
1.1.6.3 Crash Protection Log (crashprotection.log) . . . 75
1.1.7 JDBC . . . 79
1.1.7.1 JDBC Settings . . . 79
1.1.7.2 JDBC Stack Trace Filter . . . 81
1.1.7.3 JDBC Log (jdbc.log) . . . 82
1.1.8 Compression . . . 84
1.1.8.1 Compression Settings . . . 84
1.1.8.2 MIME Type Restrictions . . . 85
1.1.8.3 Exclude URLs . . . 87
1.1.9 Content Filters . . . 89
1.1.9.1 Search and Replace . . . 89
1.1.9.2 Content Filter Restrictions . . . 90
1.1.9.3 Filter Settings . . . 92
1.1.10 Plugins . . . 92
1.1.10.1 Active Plugins . . . 92
1.1.10.2 Open Services Gateway initiative (OSGi) . . . 93
1.1.10.3 Reactor Log (reactor.log) . . . 95
1.1.11 Logging . . . 95
1.1.11.1 Log Archive . . . 96
1.1.11.2 Log Settings . . . 98
1.1.11.3 Overview of FusionReactor Logs . . . 100
1.1.12 Analytics . . . 104
1.1.12.1 Connector Status . . . 104
1.1.12.2 FusionAnalytics Settings . . . 107
1.1.12.3 FusionAnalytics DataCollector Targets . . . 107
1.1.12.4 Locked . . . 110
1.1.12.5 Naming Applications and DataCollectors . . . 111
1.1.12.6 FRAM Target . . . 113
1.1.13 Fusion Reactor . . . 114
1.1.13.1 Enabling and Disable . . . 115
1.1.13.2 Notifications . . . 115
1.1.13.3 Notification Log . . . 117
1.1.13.4 Settings . . . 119
1.1.13.5 Filter Restrictions . . . 122
1.1.13.6 Change Password . . . 124
1.1.13.7 Change FR-REST Password . . . 125

1.1.13.8 Log . . . 125
1.1.13.9 About . . . 126
1.1.13.10 Help . . . 127
1.1.13.11 Logout . . . 127
1.1.14 Request Detail . . . 127
1.1.14.1 Stack Tracing . . . 127
1.1.14.2 Request Detail Page . . . 127
1.1.15 Analyzing Data . . . 128
1.1.15.1 FusionAnalytics . . . 128
1.1.15.2 Importing and Graphing Data in Excel . . . 129
1.1.16 Users . . . 129
1.1.16.1 Types Of User Account . . . 130
1.1.16.2 Managing The User Accounts . . . 131
1.1.17 Regular Expressions . . . 132
1.1.17.1 Regular Expressions and Restrictions . . . 132
1.1.17.2 Regular Expressions and Content Filters . . . 134
1.1.18 FusionReactor API (FRAPI) . . . 135
1.1.18.1 FRAPI Examples in ColdFusion . . . 135
1.1.19 FR Administration Manager (FRAM) . . . 151
1.1.19.1 Instance Manager . . . 152
1.1.19.1.1 Initial Server Scan . . . 156
1.1.19.1.2 Initial Update . . . 160
1.1.19.1.3 Add Server & Scan System . . . 161
1.1.19.1.4 Install FusionReactor Instance . . . 162
1.1.19.1.5 Uninstall FusionReactor Instance . . . 169
1.1.19.1.6 Edit FusionReactor Instance . . . 174
1.1.19.1.7 Edit FusionAnalytics Connector . . . 177
1.1.19.1.8 Update FusionReactor Instance . . . 196
1.1.19.1.9 Corrupt Configurations . . . 205
1.1.19.1.10 Unsupported Configurations . . . 206
1.1.19.2 Setting Up FRAM with HTTPS . . . 207
1.1.20 Licensing . . . 208
1.1.20.1 Uploading your License File . . . 208
1.1.20.2 Activation Period . . . 209
1.1.20.3 License Activation . . . 211
1.1.20.4 Activating Multi-Key License Files . . . 213
1.1.20.5 License Policy . . . 213
1.2 Quick Start to Troubleshooting Server Problems with FusionReactor . . . 213
1.2.1 Getting FusionReactor . . . 214
1.2.2 Configuration and Diagnosis in Emergencies . . . 214
1.2.3 Configuration and Diagnosis for Production . . . 216
1.2.4 Configuration for Enterprise and High-Availability Environments . . . 217
1.2.4.1 Setting Up A Separate Monitoring Instance . . . 218
1.2.4.2 FusionReactor in High Availability Environments (HAE) . . . 218
1.2.4.3 Enterprise Dashboard Servers and Groups . . . 218
1.2.5 Using log files to troubleshoot issues . . . 219
1.2.6 Getting Help . . . 219
1.2.7 Logs . . . 220
1.2.7.1 Stack Trace Log (stacktrace.log) . . . 220
1.3 Crash Protection and Monitoring with FusionReactor . . . 221
1.3.1 Crash Protection Basics . . . 222
1.3.1.1 Design Goals . . . 223
1.3.1.1.1 Common J2EE Problems . . . 223
1.3.1.1.2 Types of FusionReactor Crash Protection . . . 223
1.3.1.1.3 Survival Strategy . . . 223
1.3.1.2 Crash Protection Implementation . . . 224
1.3.1.2.1 Implementation . . . 224
1.3.1.2.2 Abort Strategies Explained . . . 224
1.3.1.2.3 Risks and Restrictions of Request Abort . . . 225
1.3.1.3 The Crash Protection Restrictions Engine . . . 225
1.3.1.3.1 Crash Protection Restrictions Engine . . . 226
1.3.1.3.2 Customer Demographics . . . 226
1.3.1.3.3 Restrictions Engine Rules . . . 226
1.3.2 Crash Protection In Action . . . 226
1.3.2.1 Test Pages . . . 227
1.3.2.2 Exercising Timeout Protection . . . 227
1.3.2.3 Exercising Memory Protection . . . 228
1.3.2.4 Exercising Request Quantity Protection . . . 229
1.3.2.4.1 Using the 'Queue' Strategy . . . 230
1.3.2.5 Using the Redirect Method . . . 230
1.3.2.5.1 Best Practices in Redirection . . . 231
1.3.3 Understanding Crash Protection Restrictions . . . 232
1.3.3.1 Modes . . . 232
1.3.3.2 Rule Basics . . . 232
1.3.3.3 Examples of Restrictions . . . 234
1.3.3.3.1 Excluding Batch Jobs . . . 234
1.3.3.3.2 Including Specific Hosts . . . 235

1.3.3.3.3 Including a Specific Action Page . . . 235
1.3.4 Crash Protection Alerts . . . 236
1.3.4.1 Email Alerts . . . 236
1.4 Using the FusionReactor JDBC Driver Wrapper . . . 251
1.4.1 Basic Setup of the JDBC Wrapper . . . 251
1.4.2 Gathering metrics from ColdFusion pages . . . 254
1.4.3 Pathologically bad pages . . . 255
1.5 Using FusionReactor Enterprise Scripting . . . 257
1.5.1 FusionReactor Enterprise Scripting in a Nutshell . . . 257
1.5.1.1 What Are Enterprise Scripts? . . . 257
1.5.1.2 How Do I Configure an Enterprise Script? . . . 258
1.5.1.3 When Does FusionReactor Run Enterprise Scripts? . . . 258
1.5.1.4 How Does FusionReactor Run These Scripts? . . . 258
1.5.1.5 Logging Scripts . . . 259
1.5.1.6 Operational Impacts of Scripting . . . 260
1.5.2 Using the Example Scripts . . . 261
1.6 FusionReactor 4.5.x Installation Guide . . . 262
1.6.1 Installing FusionReactor on Windows . . . 263
1.6.1.1 Installing FusionReactor 4.5.x on Windows . . . 264
1.6.1.2 Updating FusionReactor 3.x on Windows . . . 275
1.6.1.3 Updating FusionReactor 2.0 on Windows . . . 285
1.6.1.4 Updating FusionReactor 1.0 on Windows . . . 286
1.6.1.5 Console Installation on Windows . . . 286
1.6.1.6 Updating FusionReactor 4.0.x on Windows . . . 292
1.6.2 Uninstalling FusionReactor on Windows . . . 299
1.6.3 Installing FusionReactor on Linux . . . 303
1.6.3.1 Installing FusionReactor 4.5.x on Linux . . . 303
1.6.3.2 Updating FusionReactor 3.x on Linux . . . 314
1.6.3.3 Updating FusionReactor 2.0 on Linux . . . 325
1.6.3.4 Updating FusionReactor 1.0 on Linux . . . 326
1.6.3.5 Console Installation on Linux . . . 326
1.6.3.6 Updating FusionReactor 4.0.x on Linux . . . 332
1.6.4 Uninstalling FusionReactor on Linux . . . 339
1.6.5 Overview . . . 343
1.6.6 Installing FusionReactor on Solaris . . . 344
1.6.6.1 Installing FusionReactor 4.5.x on Solaris . . . 344
1.6.6.2 Updating FusionReactor 3.x on Solaris . . . 355
1.6.6.3 Updating FusionReactor 2.0 on Solaris . . . 366
1.6.6.4 Updating FusionReactor 1.0 on Solaris . . . 367
1.6.6.5 Console Installation on Solaris . . . 367
1.6.7 Installing FusionReactor on Mac . . . 374
1.6.7.1 Installing FusionReactor 4.5.x on Mac . . . 374
1.6.7.2 Updating FusionReactor 3.x on Mac . . . 385
1.6.8 Uninstalling FusionReactor on Mac . . . 396
1.6.9 Uninstalling FusionReactor on Solaris . . . 400
1.6.10 Where are FusionReactor 4's Manual Installer Packages? . . . 404
1.6.10.1 Manual Instance Installation . . . 405
1.6.11 Installing FusionReactor in Locked Down Environments . . . 408
1.7 JDBC Driver Wrapper User Guide . . . 411
1.7.1 JDBC Driver Wrapper Overview . . . 411
1.7.2 Installation in ColdFusion . . . 412
1.7.3 Creating a New Wrapped Datasource in ColdFusion . . . 415
1.7.3.1 Using Macromedia's (DataDirect) ColdFusion Built-in Drivers . . . 415
1.7.3.2 Constructing JDBC URLs for other DataDirect Drivers (CFMX,7,8,9,10) . . . 417
1.7.3.3 Using a User-Specified Driver (“other”) Driver . . . 419
1.7.4 Using the FusionReactor JDBC Wrapper . . . 421
1.7.5 Interpreting JDBC Log Data . . . 427
1.7.6 Prepared Statement Positional Bind Parameters Replacement Strings . . . 430
1.7.7 A Note On SQL Server Select Methods . . . 432
1.7.8 Exception Catalog . . . 435
1.7.9 FusionReactor 4.5.0: Upgrading to 4.5.0 from the Split-Jar Procedure . . . 436
1.8 FusionReactor Plugins . . . 437
1.8.1 AMF Dissector Plugin . . . 438
1.8.2 FusionReactor Deadlock Plugin . . . 440
1.8.3 FusionReactor Extensions for ColdFusion (FREC) . . . 443
1.8.3.1 FREC Installation Guide . . . 444
1.8.3.1.1 Installing FusionReactor Extensions for ColdFusion (FREC) . . . 445
1.8.3.1.2 Installing FREC into FusionReactor Instances . . . 449
1.8.3.1.3 Uninstalling FREC from FusionReactor Instances . . . 454
1.8.3.1.4 Uninstalling FusionReactor Extensions for ColdFusion (FREC) . . . 457
1.8.3.2 FREC Log Files . . . 461
1.8.3.2.1 CF Database Connection Statistics (dbpoolstats.log) . . . 461
1.8.3.2.2 CF Hit Count Statistics (hitcountstats.log) . . . 462
1.8.3.2.3 CF Log Entries (fr-coldfusion-log.log) . . . 464
1.8.3.2.4 CF Realtime and Request Stats (realtimestats.log) . . . 465
1.8.3.3 Configure FusionReactor ColdFusion Plugin (FREC) . . . 469
1.8.3.4 How to enable FREC . . . 470

1.8.4 FusionReactor Heartbeat Plugin 471
1.8.5 FusionReactor JRun Plugin 474
1.8.6 FusionReactor Log Rotator Plugin 475
1.8.7 FusionReactor Notification Generator Plugin 476
2. Memory Protection 478
3. Methodology in OpenOffice 479
4. Restriction 479
5. Timeout Protection and Request Protection 479

Home

Introducing FusionReactor 4.5.x

FusionReactor is a professional server monitor, designed for production environments and supporting a range of J2EE servers and Adobe servers (ColdFusion 6, 7, 8, 9 and 10 LiveCycle and BlazeDS). As a pro-active monitor, it increases server stability and performance by continuously examining applications, servers and databases to identify, highlight and address issues before they impact critical processes. FusionReactor allows you to discover, visualize and protect against common web application performance problems through graphs and logs. Whether used for production monitoring or as a developer, a tool like FusionReactor is key to solving many common server problems.

Getting Started

Downloading and Installing FusionReactor

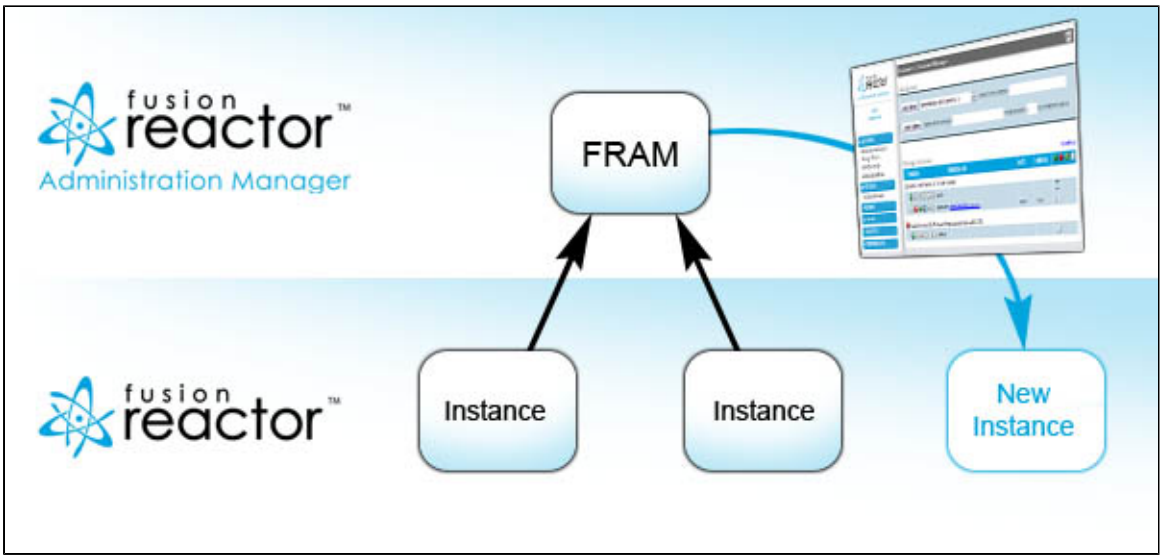
We are assuming you've downloaded and installed FusionReactor already; if not, go ahead and install it now. The installation process is easy – in almost every case the installer package will take care of all the details. Read the installation guide for details on how to install FusionReactor if you haven't already installed FusionReactor.

Licensing

After installing, you will need to upload your license and activate it. This is usually straight forward, but if you're not connected to the Internet to automatically activate your license, then click here to find out how to perform a manual activation.

Installing FusionReactor instances within FRAM

When you install FusionReactor you will automatically get an instance of the FusionReactor Administration Manager (FRAM) installed. The first time you access FRAM you will be prompted to scan for other servers into which new instances of FusionReactor can be installed.



In the above image you can see a very general outline of how FusionReactor and FRAM work together. After initially installing FusionReactor you will have FRAM installed and as mentioned above you will be given the opportunity to scan for servers. After that point you can use the Instance Manager to install new FusionReactor instances.

Once you have installed your FusionReactor instances you can use the Instance Manager to add them to the FRAM Enterprise Dashboard from where you can easily monitor the status of all configured servers. (You can also manually add instances to the Enterprise Dashboard in order to see the status of FusionReactor instances which weren't installed by the current FRAM application.)

Troubleshooting

There are many different scenarios in which FusionReactor is used, each with its own conditions, and each requiring a different configuration. The following are links to sample configurations for short term (immediate) troubleshooting and longer-term production monitoring.

- Configuring FusionReactor in emergencies
- Log files are critical to understanding what's going on, on your server - read about log file's and how to configure them.

Feedback

We welcome feedback on all our products and publications. To discover the various ways of contacting us, please use our contact page. We will address your feedback as quickly as possible.

Support Mailbox

You can mail us FusionReactor questions at:

support@fusion-reactor.com

Google Group

We also provide a Google Group for the community to ask questions in.

http://groups.google.com/group/fusionreactor

We check and reply to this group regularly, so it's also worth searching to see if your question has already been answered.

FusionReactor User's Guide

Getting Started

Requests

- Running Requests
- Request History
- Request Settings
- Request Log (request.log)
- Request Capture

Enterprise

- Enterprise Dashboard
- Manage Servers
- Manage Groups
- Enterprise Settings

- Using Monitored Server Scripts

Metrics

- System Metrics
- Longest Requests
- Slow Requests
- Longest JDBC
- Metrics Settings

Resources

- Request Activity Graph
- Request Time Graph
- JDBC Activity Graph
- JDBC Time Graph
- Memory Graphs
  - Memory Heap
  - Heap and Non Heap
  - Types of Memory Graphs
- CPU Graph
- List All Threads
- Resource Settings
- Resource Log (resource.log)
- Memory Logs
  - Memory Heap and Non Heap Log (memorysummary.log)
  - Types of Memory Logs
- CPU Thread Log (cputhread.log)

Crash Protection

- Crash Protection Restrictions
- Crash Protection Settings
- Crash Protection Log (crashprotection.log)

JDBC

- JDBC Settings
- JDBC Stack Trace Filter
- JDBC Log (jdbc.log)

Compression

- Compression Settings
- MIME Type Restrictions
- Exclude URLs

Content Filters

- Search and Replace
- Content Filter Restrictions
- Filter Settings

Plugins

- Active Plugins
- Open Services Gateway initiative (OSGi)
- Reactor Log (reactor.log)

Logging

- Log Archive
- Log Settings
- Overview of FusionReactor Logs

Analytics

- Connector Status
- FusionAnalytics Settings
- FusionAnalytics DataCollector Targets
- Locked
- Naming Applications and DataCollectors
- FRAM Target

Fusion Reactor

- Enabling and Disable
- Notifications
- Notification Log
- Settings
- Filter Restrictions
- Change Password
- Change FR-REST Password
- Log
- About
- Help
- Logout

Request Detail

- Stack Tracing
- Request Detail Page

Analyzing Data

- FusionAnalytics
- Importing and Graphing Data in Excel

Users

- Types Of User Account
- Managing The User Accounts

Regular Expressions

- Regular Expressions and Restrictions
- Regular Expressions and Content Filters

FusionReactor API (FRAPI)

- FRAPI Examples in ColdFusion

FR Administration Manager (FRAM)

- Instance Manager
  - Initial Server Scan
  - Initial Update
  - Add Server & Scan System
  - Install FusionReactor Instance
  - Uninstall FusionReactor Instance
  - Edit FusionReactor Instance
  - Edit FusionAnalytics Connector

    - Update FusionReactor Instance
    - Corrupt Configurations
    - Unsupported Configurations
- Setting Up FRAM with HTTPS

Licensing

- Uploading your License File
- Activation Period
- License Activation
- Activating Multi-Key License Files
- License Policy

Getting Started

FusionReactor is a professional server monitor, designed for production environments and supporting a range of J2EE servers and Adobe servers (ColdFusion 6, 7, 8, 9 and 10, LiveCycle and Flex Data Services). As a pro-active monitor, it increases server stability and performance by continuously examining applications, servers and databases to identify, highlight and address issues before they impact critical processes.

Once you have installed FusionReactor you will want to access the FusionReactor Administrator (FRAM).

There are two ways to do this:

- The simplest method is to use the desktop or Start Menu icons, which were created during installation.
- Alternatively you can call the FusionReactor Administrator (FRAM) directly from a web browser with a URL like this: `http://127.0.0.1:8087/fusionreactor/`

> **NOTE**
>
> If you are **using FusionReactor with Adobe ColdFusion** - we recommend that you install FusionReactor Extensions for ColdFusion (FREC)

The first page you will see when you access FusionReactor Administration Manager (FRAM) will be a login page. Enter the password fpr the ***Administrator*** user, which you entered during installation, and click on the *'Login'* button. Note that if you did not specify passwords for the Manager or Observer users, these accounts are disabled and you will be unable to login as these types of users.

Once you are logged in, you will notice that you have a 10 day trial to use FusionReactor. Please check the license guide for a detailed guide on how to acquire and activate your license.

Instance Manager

The Instance Manager is used to keep track of the FusionReactor instances you may have monitoring your servers. Once you have installed a FusionReactor instance into a server instance, you are able to add that instance to the Enterprise Dashboard where you have an overview page of different metrics (The enterprise dashboard is only available with enterprise licenses).

If you start FRAM from the installer, the Initial Server Scan wizard will pop up and the FusionReactor Administration Manager will attempt to scan for supported ColdFusion and J2EE servers. For a detailed guide on this process, please see Initial Server Scan. If you go through the shortcut or direct URL, the initial server scan wizard will not pop up and you will have to initiate the server scan yourself on the Instance Manager.

You are able to monitor any number of installed ColdFusion and J2EE servers. To install a FusionReactor Instance, click the *'Add Instance'* button next to the desired server instance you want to monitor. Follow the Install FusionReactor Instance guide for more detailed information.

Logging

The logging functionality in FusionReactor is very powerful. Lots of the data shown in the metrics within FusionReactor instances is stored in-memory so it can be used to generate graphs and reports. If a server was to crash or be shutdown, this in-memory data is lost - however all of this data will have been captured to one of FusionReactor's extensive log files where you will be able to determine what went wrong.

On the Logging pages, you are able to set up log rotation to send the logs to a FusionAnalytics Data Collector, where you can have FusionAnalytics import the application log files into a SQL database. The resulting data is processed by the FusionAnalytics Server and visualized by the FusionAnalytics Desktop application. To be able to do this, you must add a FusionAnalytics target.

For a more detailed overview on FusionReactor's logging, please go the Overview of FusionReactor Logs guide.

FusionAnalytics Connector

The Analytics menu provides access to the current status between your instance of FusionReactor and FusionAnalytics. FusionAnalytics is a ColdFusion Application and server analysis tool, which integrates seamlessly with FusionReactor. FusionAnalytics is all about "Making IT Better" and giving you the information and data to help you make better business decisions, improve application performance and quality of your applications as well as measure exactly how your applications are performing on a continuous day to day basis.

The FusionAnalytics Connector Status page will show details such as the size of the archive, the source of the archive (which instance it came from), any exceptions and what target the archive has been set. You can also see all of the archive files (which include the log files) that are going to be transferred to FusionAnalytics. You are able to rotate to have FRAM generate new archives based on the log information gathered since the previous archives was generated.

The archived files are sorted by date, making finding archives easy. You can also delete archives, where you can either delete them individually or you can delete all of the stored archived in one go by clicking on '*Clear all Archives*'

> **WARNING**
>
> Make sure you are 100% sure before removing all archives as they will be deleted from your hard drive.

For a more detailed overview, please go to the Connector Status page.

Requests

Running Requests

Request History

Request Settings

Request Log (request.log)

Request Capture

Running Requests

Running Requests give you the ability to view and administer running requests on your server. From the FusionReactor Administrator, click Requests > Running Requests and you will see the Running Requests screen. At the top of the screen you will see the refresh bar. You can use this to manually or automatically update the running requests listed below. For each request shown you will see the time at which it began, the IP

address of the machine which created it, the page which was actually requested, the amount of time it has been running and the state of the memory when this request entered the system. Above the running requests you will also see the state of the memory right now. The request table has the following columns:

Name	Description
Started	The time at which the request was received.
IP	The IP address of the caller.
ID	The request ID and the name of the thread which is servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request has been running so far.
Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown:Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit). The Running Requests page lists all currently running pages as well as showing the current memory values.
Max	The amount of memory available to the JVM (this is a hard upper limit).
Alloc	The amount of memory currently allocated by the JVM for its heap – this may grow until the Max value is reached.
Free	The amount of free memory from the currently-allocated heap.
Used	The amount of memory which is actually in use. This is expressed as a percentage as well as in KB.

Queued Requests (due to Crash Protection)

If there are requests queued due to Crash Protection then they will be listed below the currently running requests. The queued requests are listed in the same format as the running requests (see above) but there is an extra column giving you the reason that the request was queued.

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles. Clicking on a title for a second time will reverse the order on that column.

Managing Requests

Each request listed on this page has four icons next to it. They are:

Button	Name	Description
	Stack Trace	Clicking this button will take you to the Stack Trace page, allowing you to see the current Java stack trace of the request.
	Request Detail	Clicking this button will take you to the Request Details page, allowing you to access all information associated with a request; e.g cookies, JDBC, headers, User Trace Markers and AMF information.
	Kill Request	Clicking this button allows you to manually kill any tracked request. You will be asked if you are sure. Click on OK to kill the request. The system will try for up to 60 seconds to kill the request (the "Soft Kill" phase) cleanly. If the request is not producing output, this method may be unsuccessful. You may also click "**Force Thread Death**" to perform a hard kill action on the request. If "Stop Tracking This Request" is checked, the request will be exempt from metric tracking - see "Stop Tracking" below.
	Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.

	Stop Tracking	Clicking this button will cause FusionReactor to stop tracking this request for certain metrics (slowest requests, request runtime. request activity). Any JDBC statements currently running will be logged and visible in Request Details, but the JDBC Time and Activity graphs, as well as Average DB Time, will not reflect them. Once this button has been pressed, it is grayed and disabled. There is no way to reverse this operation.

There are two overall options:

Button	Description
Stack Trace ALL	Displays a stack trace of all currently-running requests. The stack trace is colorized to ease recognition of each part.
Kill ALL	Allows the user to kill all currently-running requests. There is an intermediate screen which allows the user to select whether the request are to be excluded from tracking (default) or not.

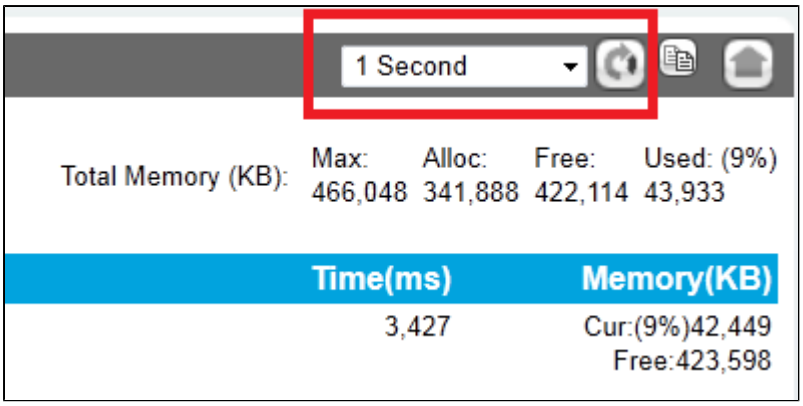
Warning - Kill Instability

The "Kill ALL", and "Kill Request" buttons do not give requests the opportunity to complete in their own time. Using this feature *can* theoretically cause your application to become unstable, because all locks are immediately relinquished by the thread when it's killed. Any data protected by these locks will retain the state it had when the thread was killed.

You might consider this a small price to pay to keep your server up though!

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh. Below is a screen shot showing this.



Request History

The Request History page shows a list of completed requests. Please be aware that FusionReactor only keeps a fixed amount of requests in memory. If the history size gets too long or you restart you application server then old requests will be removed. To change the size of the request history, change the History Size value on the Request Settings page. Note that you can still examine these older requests by looking into the Request Log (request.log).

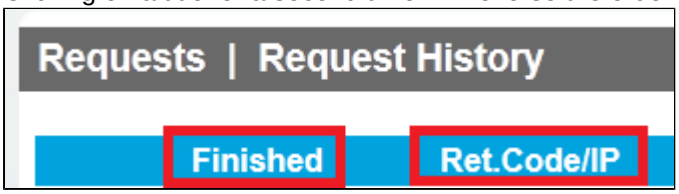
The request history table has the following columns:

Name	Description
Finished	The time at which the request was completed.
Ret.Code/IP	The final HTTP return code of the request and the IP address of the caller. If the request was killed or if it triggered Crash Protection then you will see a note here letting you know what happened.
ID	The request ID and the name of the thread which is or was servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request took to complete. If available then you will also see the amount of actual CPU time that this request used below it.
Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown: Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit).

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles.

Clicking on a title for a second time will reverse the order on that column. Below is a screen shot showing this.



Managing Requests

Each request listed on this page has two icons next to it. They are:

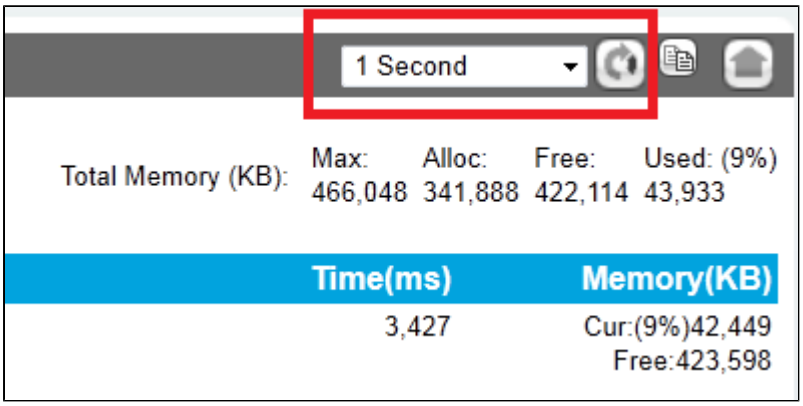
Button	Name	Description
	Request Detail	If you click this button then you will be taken to the Request Details page which gives you access to all information associated with a request. eg. Cookies, JDBC, Headers, User Trace Markers and AMF data.
	Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request

Clear

The Request History list can be cleared completely by clicking the **Clear History** button.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh. Below is a screen shot showing this.



Request Settings

The Request Settings page lets you modify various elements relating to the capture, storage and display of Request information. The page has the following items:

Request Formatting

Configuration	Description
URL Format	This option changes how URLs are displayed throughout FusionReactor. "Full" includes the hostname, while "Short" just shows the path. Select "Full" if you are dealing with servers which host multiple websites.

Request Logging

Configuration	Description
Request Logging	If Request Logging is enabled then requests will be written to the request.log as they arrive and complete. If you disable it then it won't.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the Request Log from inside FusionReactor.

Request Content Capture

Configuration	Description
Request Capture	Activating this option instructs FusionReactor to store the complete Request and Response objects to disk. **WARNING** Turning on Request Content Capture will create multiple XML files for every request run on a server. This may use very large quantities of disk space quickly. Please be careful when enabling this feature and do not leave it enabled if you are not monitoring it.
Capture Path	The location where captured requests and responses are stored.

History

Configuration	Description
History Size	This value defines how many requests FusionReactor stores in memory. Once the Request History reaches said number, old requests will be removed when new ones arrive.

Stream Metrics

Configuration	Description
Stream Metrics	Specifies whether FusionReactor should track statistics about how long it took to get a response back to the client. (Time to first byte, time to completion, bytes per second, etc.).

Request Log (request.log)

The Request Logs (a set of rotating log files accessible over the Request Settings page) are simple space-separated files so they can be easily imported into database or spreadsheet applications. Note that the amount of data stored in the log files is typically far greater than the number of requests shown on either the Request History or Request Log screens. After a server restarts the Request Log screen will show the last entries on the Request Log, even if these requests were run before the server restart. You can change the amount of data shown on the Request Log screen by changing the View Size on the Request Settings page. Refer to Log Settings to see where the logs are being stored.

The log file is formatted as follows:

Field	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The date on which this log entry was written.
Date/Time (ms)	3 [C]	This is the millisecond time at which this request was started or finished (depending upon the Request Status.)
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
Request ID	6 [F]	The FusionReactor Request ID.
Request Status	7 [G]	The current state of the request. This will tell you if a request has started, finished, was killed, queued, etc...
CP Reason	8 [H]	The reason given by Crash Protection for the current action.
Thread ID	9 [I]	The name of the thread responsible for responding to this request.

Client IP Address	10 [J]	The IP address of the machine making the request.
Request Method	11 [K]	This will usually be "GET" or "POST"
Request URL	12 [L]	This is the requested URL. The format of this column will change depending on URL Format key which can be found on the Request Settings page.
Execution Time (ms)	13 [M]	The amount of milliseconds it took to complete the request. (For incomplete requests, this column will be 0.)
Used Memory (percentage)	14 [N]	The amount of memory (expressed as a percentage) which was used when this request started or finished.
Max Memory (KB)	15 [O]	The total amount of physical memory available to this instance.
Used Memory (KB)	16 [P]	The amount of memory which was used when this request started or finished.
Total Memory (KB)	17 [Q]	The amount of memory which was allocated by the instance when this request started or finished.
Free Memory (KB)	18 [R]	The amount of free memory (within the allocated block) when this request started or finished.
Query String	19 [S]	If the URL has any parameters then they will appear here.
Return Status Code	20 [T]	This is a HTTP return code such as 200 (OK,) 404 (Not found,) or 500 (Internal Server Error.) For "Started:" rows this column will be 200.
CPU Time (ms)	21 [U]	This is the amount of actual CPU time which this request required. (For incomplete requests, this column will be 0.)

AMF Request	22 [V]	If AMF decoding is enabled (Request Settings) and there is AMF to decode, then the method names will appear here.
JSESSIONID	23 [W]	The J2EE Session Id for this request.
CFID	24 [X]	The CF Id for this request.
CFTOKEN	25 [Y]	The CFTOKEN for this request.
JDBC Query Count	26 [Z]	The number of JDBC queries run by this request at the time of this log.
JDBC Total Time (ms)	27 [AA]	The total amount of time spent running JDBC queries at the time of this log.
JDBC Total Execution Time (ms)	28 [AB]	The total amount of time spent by the database running JDBC queries at the time of this log.
JDBC Total Rows	29 [AC]	The total number of rows returned by JDBC queries at the time of this log.
Bytes Sent	30 [AD)	The amount of data which was sent back to the client.
Time to First Byte (ms)	31 [AE]	The number of milliseconds it took to deliver the first bit of data.
Time to Last Byte (ms)	32 [AF]	The number of milliseconds it took to deliver the complete content.
Time to Stream Open (ms)	33 [AG]	The number of milliseconds before the data stream was opened.
Time to Stream Close (ms)	34 [AH]	The number of milliseconds until the completed data stream was closed.

Deleting the Request Log

Clicking on the Delete Log File button will display an "Are you sure?" dialog box. If you OK this dialog then the Request Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See:

Log Settings
Importing and Graphing Data in Excel

Request Capture

Request Capture is the ability to store incoming requests and their associated responses to the file system. When request capture is set to requests and responses you will end up with three files for each request. A request file, a body file and a response file. The header XML file holds all information sent to server by the client. The response XML file holds everything sent back (besides body) and the body file holds the actual returned content.

> **WARNING**
>
> Turning on Request Content Capture will create multiple XML files for every request run on a server. This may use very large quantities of disk space quickly. Please be careful when enabling this feature and do not leave it enabled if you are not monitoring it.

How do I set up Request Capture?

From the FusionReactor Administrator, click Request Settings from the table of contents (Requests section) and you will see the Request Settings page. Scroll down until you reach a section named "Request Content Capture". From here you can control the request capture feature. Request Capture can be in three states. Disabled, Capture Requests only, or capture both requests and responses. The Capture Path is the directory into which the captured requests will be stored.

Why would I want to Capture Requests?

Request capture can be useful for a variety of reasons from technical to legal, these include:

Analysis

The FusionReactor stored requests hold much more information than you would normally expect to find in web logs and as such they have a much higher value for analyzing your application traffic.

Development

You could use stored requests in order to be able to reproduce specific load conditions or test cases.

Support

Anyone who has had to have a user describe the contents of their screen will instantly realize how much easier life would be if you had an exact copy of their response in easy access.

Enterprise

Enterprise Dashboard

Manage Servers

Manage Groups

Enterprise Settings

Using Monitored Server Scripts

Enterprise Dashboard

The Enterprise Dashboard lets you track the state of installations on one or more servers. When you install an Enterprise version of FusionReactor on a server it can be queried for its current state. The Enterprise Dashboard lets you instantly monitor the health of ALL of your server instances/machines from a single window. You immediately see the state of memory, CPU, DB load and average request time, in an easy to read graphical dashboard display.

By default you will see your install server in the Enterprise Dashboard but, using the Manage Servers page, you can add any other server which is running FusionReactor.

Once you have your machines set up, go to the Dashboard. The default view when bringing up the Enterprise Dashboard is to show all groups (which can be set up on the Manage Groups page) along with any servers which are currently ungrouped.

Group Icons

You can group machines together to help organize your dashboard, see Manage Groups for more information. Each machine and group is represented by an icon which will change color to indicate its current health. As well as this, clicking on a server or group will give you its most recent statistics. The dashboard also provides links to each monitored instance so it is a simple job to investigate a server further.

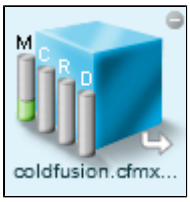
Clicking on a group icon will display all servers within that group. You will also see the Group Information Pane which will display statistics about the group and the servers within. By default the group icon will be displayed in blue, but if any grouped server is having problems then the color of the icon will change and a small warning panel will be displayed to let you know what the problem is.

How Do I Add Machines?

There are three ways to add servers to your Enterprise Dashboard: Firstly, you can import servers from your Instance Manager. Secondly, you can get FusionReactor to scan its machine for servers, or you can simply insert the URL of your server into the Manage Servers page.

Server Icons

The server icon shows four bars. These bars each have a letter above them to indicate what is being displayed.



These bars are:

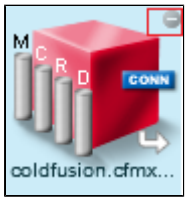
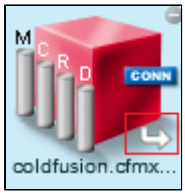
Bar	Name	Description
M	Memory	This represents the amount of memory currently being used by this server.
C	CPU	This represents the amount of CPU time used by this instance since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.

R	Request Time	This represents the average request time on this server over the last 60 seconds.
D	DB Time	This represents the average time spent on DB operations in the last 60 seconds.

Clicking on a server icon will display the Server Information Pane which will display statistics about the server. By default the server icon will be displayed in blue, but if there are any problems with the server then the color of the icon will change and a small warning panel will be displayed to let you know what the problem is.

Server Icon Buttons

Additionally, there are two small buttons on the Server icon:

Button	Name	Description
	Stop monitoring	At the top right corner of the icon is a small button which can be used to flip a server between Online and Offline status. An offline server will be displayed in gray and will not cause warnings to appear within the Enterprise Dashboard.
	Open FusionReactor in a new Window	At the bottom right corner of the icon is a small arrow. If you click on this arrow then you will be taken directly to the instance of FusionReactor running on this server.

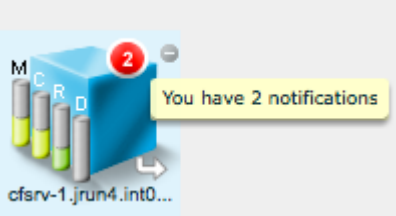
Warning Lights

When FusionReactor detects a problem on a monitored server it will light up one or more "Warning Lights" on the server cube. You can move your mouse over a warning light to see a tool tip explaining the problem is in more detail. The following table describes the warning lights and their meanings:

Message	Meaning	Description
CONN	Connection Lost	The connection (heartbeat) has been lost to the monitored server. The monitored server may be down.
SLOW	Slow Requests	The number of requests running longer than the slow page threshold (within the last 60 seconds) has exceeded the Warn or Critical values configured on the Enterprise Settings page.

LONG	Long Running Requests	The average request time (over the past 60 seconds) has exceeded the Warn or Critical value configured on the Enterprise Settings page.
MEM	Low Memory	The used memory percentage exceeded the Warn or Critical value configured on the Enterprise Settings page.
QTY	CP Request Protection	The number of triggered crash protections (for the request quantity limit) exceeded the Warn or Critical value configured on the Enterprise Settings page.
TIME	CP Timeout Protection	The number of triggered crash protections (for the request timeout limit) exceeded the Warn or Critical value configured on the Enterprise Settings page.
DB	JDBC Slow Requests	The average JDBC request time (over the past 60 seconds) has exceeded the Warn or Critical value configured on the Enterprise Settings page
ACT	Please Activate License	You need to activate the license installed on the monitored server.

Notifications



The number inside the red notification circle (to the top right of the server icon) represents the number of current Notifications (messages from FusionReactor's various subsystems that alert you to conditions which require attention, or are otherwise abnormal) the server currently has. To see the notifications in more detail, simply click the number to go to the Notifications page.

See here for information on Notifications

Group Information Pane

The group information pane shows you how many servers are within this group as well as showing statistics for all grouped servers. These statistics are:

Name	Description

Group Size	This shows the number of servers within this group.
Request Count	This shows you how many requests are running on all grouped servers along with the total number of requests which have been run.
Recent Slow Pages	This shows how many requests within the last 60 seconds have run longer than the Slow Page Threshold value.
Average Request Time	This is the average request time for all grouped servers for the last 60 seconds, and for the complete history size.
Average DB Time	This is the average time spent on database operations for all grouped servers.
Average Memory	This is the average used and total allocated memory for all grouped servers.

Additionally, you will see two graphs (Average Used Memory and Request Count.) These graphs show the same information displayed in the statistics section but they are presented visually for ease of use.

Server Information Pane

The server information pane shows you various statistics about your selected server. Statistics which are listed with blue titles are links. Click on one of these titles to be taken to a more detailed page within the instance of FusionReactor running on that server. If a server is currently having problems then the appropriate statistic will be highlighted to indicate where the problem is. The displayed statistics are:

Name	Description
FusionReactor Instance	This is the full ID of the instance of FusionReactor running on this server.
License	This shows the type of FusionReactor license running on this server.
Activation Status	This shows you if the license on this machine has been activated or not.

Crash Protection (Page Aborts)

Name	Description
Request Quantity	This shows the number of pages which triggered Crash Protection due to request quantity in the last 60 seconds and in total.
Runtime Timeout	This shows the number of pages which triggered Crash Protection due to runtime timeout in the last 60 seconds and in total.

Low Memory	This shows the number of pages which triggered Crash Protection due to low memory in the last 60 seconds and in total.

Overview

Name	Description
Server Uptime	This shows how long this server has been up.
Current Request Count	This is the number of requests currently running on this server.
Total Request Count	This is the total number of requests which have run on this server.
Average Request Time	This shows the average request time on this server over the last 60 seconds, and for the complete duration of it's request history.
Used Memory	This is the amount of memory currently used by this server (not the amount currently allocated.)
Allocated Memory	This is the amount of memory that has been allocated to the server.
Maximum Memory	This is the total memory that the server has.
Average Instance CPU	This shows the amount of CPU time used by this instance since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.
Average DB Time	This shows the amount of time spent on database operations on this server.
Queue Size	This is the number of requests currently queued due to Crash Protection

Recent

Name	Description
Recent Slow Pages	This shows how many requests within the last 60 seconds have run longer than the Slow Page Threshold value.

Additionally, you will see seven graphs (Req Activity, Av Request Time, Req Count, JDBC Activity, JDBC Time, Memory and CPU). These graphs show the same information displayed in the statistics section but they are

presented visually for ease of use. All graphs (except Req Count) have an arrow on the top right of them. Clicking this will take you to a more detailed graph with additional functionality (time lines) within the instance of FusionReactor which is running on this server.

Configuring the Enterprise Dashboard

Many elements within the Enterprise Dashboard can be modified. In order to make changes, please refer to the Enterprise Settings page. From there you can:

- Specify whether you should receive an email when servers start up and shut down.
- Specify what the maximum values should be on graphs and bars.
- Specify the warning and critical state threshold values for various statistics (Warn and Critical).

Manage Servers

From this screen you can add new servers to the Enterprise Console or manage existing servers. It is also possible to directly add local instances to the Enterprise Dashboard from the Instance Manager.

The form to add a new server has the following fields:

Field Name	Description
Name	This is the name which will appear in the Enterprise Dashboard under your server icon.
URL	This is the location of FusionReactor on the target server. An example URL will look something like this: http://127.0.0.1:8088/fusionreactor
Password	This is the password of FusionReactor on the target server.
Script	If you specify the path to a script on your server then it will be called when the server goes down or comes up. See the script section below for more information or check out Using FusionReactor Enterprise Scripting
Status	You can choose if this new server is online or offline. An online server will collect statistics and display them in the Enterprise Dashboard and if there are any problems then you will see warnings. If a machine is offline then no information will be gathered and the icon will be displayed in gray. If you are adding a server which is currently turned off then you should set it's status to offline. You can change it's status later by returning to the Manage Servers page (clicking the green button) or by clicking the status button (top right hand corner of a server icon) directly within the Enterprise Dashboard.
Groups	If you have any groups set up then they will be listed here. Simply select which groups you want this new server to appear within. You can add new groups at any time from the Manage Groups page.

Once you submit the form the page will be refreshed and your new server will be listed at the bottom of the page along with its current connect status.

Managing Existing Servers

Below the Add Server form you will see a list of all servers currently available in the Enterprise Dashboard. Each server has three icons next to it:

Button	Name	Description
	Modify Server	Clicking on this icon will take you to the Modify Server page. From this page you can change any of the information which is listed above.
	Check Server	Clicking on this icon will check the connect status for a server. The screen will be refreshed and you will see a message telling you if there were any problems trying to access FusionReactor on this server.
	Delete Server	Clicking on this icon will prompt you if you are sure you want to delete the server. If you OK the dialog then the server will be removed from the Existing Servers list and from the Enterprise Dashboard.

Clicking on the URL for any server in the Existing Servers list will take you directly to the instance of FusionReactor which is running on that server.

Shutdown / Startup Script

The script file which you can associate with a server will get called whenever that server goes down or comes back up again. It takes 5 parameters, they are:

Parameter	Description
Status	This parameter will have the value "DOWN" or "UP which correlates to the status of the server.
InstanceName	This is the instance name of the server in question (as defined on the Manage Servers page.)
InstanceIP	This is the IP address of the server in question.

PID	The Process ID of the Server. You could use this in order to terminate a server which had stopped responding. (If the PID is unknown then a value of -1 will be passed.)
LastSeen (ms)	This is the last time (expressed in milliseconds) that the server was last successfully contacted.

Manage Groups

From this screen you can create or modify groups in order to better organize your servers within the Enterprise Dashboard. The Add Group form is very simple:

Field Name	Description
Name	This is the name of your new group which will appear in the Enterprise Dashboard and on the Manager Servers screen.
Servers	If you have any servers currently set up then they will be listed here. Simply select which servers should be added to your new group.

Managing Existing Groups

Below the Add Group form you will see a list of all existing groups. Each group has two icons next to it:

Button	Name	Description
	Modify Group	Click this icon to change a group name or to change the servers which are currently a part of this group.
	Delete Group	Clicking on this icon will prompt you if you are sure. OK the dialog box to delete the group and refresh the page.

Enterprise Settings

The Enterprise Settings page lets you modify various elements relating to the Enterprise Dashboard. The page has the following items:

Server Shutdown/Start Up Alerts:

Configuration	Description

Send Alert Email	With this option you can decide if you want to be sent a notifications (via email) when a machine starts up and shuts down, just shuts down or you can opt to receive no notifications (Disable option). To receive notifications you must first configure your email server on the FusionReactor Settings page.

Heartbeat Settings for Monitored Instances:

Configuration	Description
Heartbeat Metrics Fetch Interval (seconds)	This interval tells FusionReactor how long to wait between each fetch of a monitored instance's metric data (its heartbeat). The heartbeat contains the data that is shown on the Enterprise Dashboard. **NOTE** Increasing this value will lower the refresh rate of the metric data shown on the Enterprise Dashboard. Decreasing this value may cause heartbeat failures. It is NOT recommended to set this value below 5 seconds.
Heartbeat Timeout (seconds)	This timeout governs how long FusionReactor will wait for a heartbeat from a monitored server instance. If no heartbeat is received within this time, FusionReactor will count it as a heartbeat failure.
Heartbeat Failure Threshold	This setting tells FusionReactor how many times the heartbeat may fail on a monitored server instance before considering it to be DOWN. FusionReactor will then report the server instance as [CONN] (Connection lost) on the Enterprise Dashboard and will run the self-healing script for the instance (if configured). If you are receiving monitoring emails and log messages for instances known to be available, try increasing this value. A value of zero will cause the first missing heartbeat to signal the server instance is DOWN.

Request Count:

Configuration	Description

Graph Y Scale	This number will be used as the maximum value on the request count graphs within the Enterprise Dashboard. **NOTE** This value is only used for display purposes. Entering a number here will NOT limit the number of requests a server will accept.

Memory:

Configuration	Description
Memory Status (percent)	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the used memory percentage exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.

Average Request Time:

Configuration	Description
Average Request Time (ms)	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the average request time (over the past 60 seconds) exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.
Graph Y Scale (ms)	This number will be used as the maximum value on the average request time graphs and bars within the Enterprise Dashboard. **NOTE** This value is only used for display purposes. Entering a number here will NOT limit the amount of time requests can run on a server.

JDBC:

Configuration	Description
JDBC Status (ms)	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the average DB time exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.
Graph Y Scale (ms)	This number will be used as the maximum value on the average DB time graphs and bars within the Enterprise Dashboard. **NOTE** This value is only used for display purposes. Entering a number here will NOT limit the amount of time database requests can run on a server.

Crash Protection:

Configuration	Description
CP Status	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the number of triggered crash protections (for a given type: request quantity, request timeout or low memory) exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.

Slow Requests:

Configuration	Description

Slow Page Status	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the number of requests running longer than the slow page threshold within the last 60 seconds exceeds then the Warn value will appear in orange. If it exceeds the Critical value then it will appear in red. You can change the slow page threshold from the Metrics Setting page.

Using Monitored Server Scripts

FusionReactor Enterprise Scripting enhances FusionReactor's Enterprise Monitoring to trigger configurable self-healing scripts when a server's responsiveness status changes.

This feature can be used to perform any task which can be launched from a platform script.

Examples might include:

- Interacting with SNMP or enterprise monitoring/reporting systems
- Performing automated restarts of affected instances
- Sending custom email messages
- Writing custom log messages

Since FusionReactor Enterprise Scripting is able to launch any platform executable (shell script,executable binary etc.), it may be tailored to virtually any environment.

> **NOTE**
>
> To use this functionality **AT LEAST** two different instances of FusionReactor Enterprise are **REQUIRED**. It is important to understand this limitation of scripting. A script should not be configured to run against the same server instance that FusionReactor is running in. If the server instance fails then the FusionReactor running in this server instance may not be able to run the script, meaning it may not be able to restart itself for example.

What Are Enterprise Scripts?

FusionReactor (Enterprise Edition) can trigger a script when a server becomes unresponsive, and when it becomes responsive again. This mechanism might allow you to pro-actively attempt to restart a failed server or instance, integrate FusionReactor into an existing monitoring environment or provide custom logging or reporting. Any program which is runnable on your FusionReactor monitoring system may be used by an Enterprise Script target.

Enterprise scripts on Windows systems, are anything that is runnable as a binary, or can be run from the command prompt, for instance:

- Binary programs
- Batch files

Enterprise scripts on Unix and Unix-based systems, this includes everything you can start from a command line, for instance:

- Binary programs
- Shell scripts (including Bash, Ruby, Python and Perl)

- Java programs (when launched from an appropriate shell script)

How Do I Configure an Enterprise Script?

Enterprise scripts are configured by editing the Script of the server's property page, accessible in the Enterprise -> Manage Servers page, then clicking on the Modify icon of the required server. You should take care to ensure the full path and filename to the script are correct. Also note that it is recommended to only configure one script for a server. If the same server has multiple scripts configured (for example you have entered the server into the dashboard in multiple groups) or you monitor the same server from different machines, then more than one script may be launched at the same time if the server has problems. The scripts may interfere with each other especially if they are trying to restart the same instance.

When Does FusionReactor Run Enterprise Scripts?

Enterprise Scripts are run whenever the Enterprise Monitor detects that a monitored instance has changed state:

- an instance which was previously available is no longer providing Enterprise data
- an instance which was previously unavailable has begun to provide Enterprise data

Scripts are run only if an instance changes state while it is being observed. Additionally, scripts are only run if:

- The Enterprise Server Alerting system is running
    - This is configured in Enterprise Settings -> Server Shutdown/Start Up Alerts
- At least "On Shutdown" must be selected.
    - If you wish to run scripts when an instance becomes available, "On Shutdown and Start Up" must be selected.
    - If you do not wish to additionally receive email for these events, disable notification in FusionReactor -> Settings -> Email Server -> Notification.

How Does FusionReactor Run These Scripts?

Launch Mechanism

FusionReactor runs these scripts by spawning them using Java system commands. The scripts will be run in the context of the user under which your J2EE (ColdFusion) server runs. This user must have at least **read + execute** access to these scripts. Any files or other executables called by the script must also be accessible by this user.

The script will be run with the current working directory (CWD) of the J2EE application server. Because of the variety of platforms available, this may be unpredictable. Any scripts you write should therefore not use the current directory notation (usually a single dot) to address files. If you plan to access files within the script, their paths should be specified completely.

Script Arguments

FusionReactor supplies several command-line arguments to the script. These arguments may be used by the script to perform logging or restart operations. Note that these arguments are supplied by FusionReactor and you do not need to enter them in the script field on the Managed Servers page. The supplied arguments are (in order):

Script Parameter	Description
UP or DOWN	reflecting the instance status
instance name	as registered in the Manage Servers screen

IP Address	as returned from a DNS lookup of the machine name part of the URL used to monitor this instance
Process ID	If available, the process ID of the J2EE application server on the remote machine. If the FusionReactor native library is not available, or FusionReactor could not read this value, this field will be -1
Last Seen Time	The time, measured in milliseconds from midnight on January 1st 1970, which the server was last successfully polled for Enterprise data. If the server has not been observed as running during this session, this field will be -1

Logging

When FusionReactor fires a script, an appropriate message is written to the Crash Protection log, located in FusionReactor/instance/<instance_name>/log/crashproteciton-0.log. This log is shared with other Crash Protection messages, and not all fields are used by Enterprise scripting.

For the exceptional cases **SCRIPTREADFAILED** and **SCRIPTEXCEPTION**, FusionReactor will log the message associated with the exception to the FusionReactor log.

Operational Impacts of Scripting

There are a few points which you should bear in mind when configuring scripting.

System Restarts and Self-Monitoring

If FusionReactor is configured to monitor itself, i.e. it is monitoring the same instance in which it is configured, scripting should not be used for operations.

In these circumstances, we recommend transitioning your environment to a High Availability monitoring solution. This entails installing a new J2EE server (Tomcat, for instance), and installing FusionReactor into that. This container will be used purely as a FusionReactor host, and will be used to monitor other containers. It may be necessary to create scripts which perform system reboots. Again, we recommend a careful evaluation of the impacts of this type of script before implementation. A script which restarts a system should not attempt to restart the system on which the monitoring solution runs.

Manual Restarts

If a script is configured for a given instance, it will be fired when that instance becomes unavailable. FusionReactor does not differentiate between overloaded (or failing) instances, and instances which have been deliberately stopped. Therefore, if you stop an instance manually, through Windows' Services panel for instance, FusionReactor will fire the configured script. As an operational matter, the affected instances should be "**offlined** " from FusionReactor before being shut down. This can be done within the Enterprise Dashboard, by clicking the (+/-) button on the top right of server icon or from within Manage Servers by selecting the Modify icon for the affected server, then changing its Status to Offline. FusionReactor will not monitor these systems. When maintenance is complete, the servers should be "**onlined**" again by reversing the process.

Using the Example Scripts

We have provided several restart scripts to get you started. This section will help you understand how to install and configure them.

Installation

The example scripts are provided in /FusionReactor/etc/cp/, thereafter the structure is split into scripts which will

run on Unix platforms, and those which will run on Windows platforms. You are free to run these scripts from this location, but we would recommend you copy these templates before editing them. You will then always have a pristine copy available for new scripts.

Worked Example: Controlling Windows ColdFusion 8 from Windows

In order to get you started, we've provided you with a worked example. In our example scenario, we will use a FusionReactor Enterprise Edition instance on a ColdFusion 7 instance to monitor a ColdFusion 8 installation, also on Windows and running FusionReactor Enterprise Edition.

Enterprise Dashboard

The first stage in preparing the environments is to ensure that both systems are running smoothly, and the monitor is able to poll the target system for enterprise data. Add the remote system to the monitor and check that Enterprise Dashboard is retrieving information from it.

Script Preparation

For this example, we'll be using the restart-Coldfusion8-OnWindows.bat script from the FusionReactor/etc/cp/windows folder. For our example, we copy the example script to a temporary folder, from where we can work on it:

copy restart-Coldfusion8-OnWindows.bat c:\tmp

In order to customize the script, we open it in an editor. All provided example scripts are commented extensively.

There are a couple of variables we must customize in the script:

- We set the LOGFILE (line 43) to c:\tmp\script.log
- We change the USER and PWD (lines 53 and 54 respectively) to reflect the Windows user with permissions to restart ColdFusion.

Adding the Script to Manage Servers

The final step in the configuration is to add the script to the monitored server's configuration. We edit the server's Enterprise Dashboard configuration by clicking on Manage Servers, then clicking the edit icon of the monitored server. We enter the script location in the Script field.

Testing the Script

The script can be tested by simply using the Windows Service control panel to stop the monitored ColdFusion 8 service. Observing the script log file c:\tmp\script.log file shows the output of the script. The ColdFusion 8 service can be observed restarting in the control panel.

Conclusion

We've shown you how to configure Enterprise Scripting to restart a ColdFusion 8 server. The scope for what scripts can do is immense, since there are no restrictions on what they may call. It would be a simple task, for example, to integrate FusionReactor into an SNMP monitoring solution, write custom log messages or send SMS text messages.

Local Monitoring: Monitoring Instances on the same computer as the Enterprise Monitor

Running all FusionReactor instances on the same computer makes it easy to develop scripts for use with Crash Protection.

Simple Watchdog

Two instances run on the same computer. One of the instances acts as an Enterprise Monitor (Watchdog) for

the other, monitored server. The watchdog server has added the monitored server to the list of managed servers in the Enterprise Dashboard of FusionReactor. A script for the monitored server is configured that will be executed by the watchdog server if the monitored server is unavailable or becomes available again. The script is used to restart the monitored server automatically after it became unavailable.

However, if the watchdog server itself becomes unavailable, the server is not longer monitored and can not be restarted automatically any more. To get around this the monitored server could also monitor the watchdog server which is described in the next section.

Cross Monitoring

One instance/server is created for the task of Enterprise Monitor to monitor the other Operational servers. Every operational server is added to the Enterprise Monitor's list of managed servers in the Enterprise Dashboard of FusionReactor. For every server a script is configured that will be executed when the server becomes unavailable (or available again). The script is used to restart the operational server automatically if it becomes unavailable. One of the operational servers is also configured to monitor the Enterprise Monitor, and a script is configured to restart the Enterprise Monitor should it become unavailable. If this way the Enterprise Monitor watches over all of the operational servers, and one of the operational Servers watches other the Enterprise Monitor. Alternatively a second Enterprise monitor could be configured with the task of watching over the first Enterprise monitor if you do not want operational servers to perform monitoring tasks.

Distributed Monitoring

Running, monitoring and restarting FusionReactor instances in a distributed environment requires remote connections between the different machines. If a monitored server becomes unavailable this is monitored by a different machine on the network. This machine then calls a script which has to connect to the remote machine and restart the remote server/instance. Depending on the operating system the participating machines use, this can be done with (SSH) or some similar technology.

Simple Watchdog

Two servers run on different machines that have a network connection. One of the servers acts as the Enterprise Monitor (watchdog) for the other, monitored server. The watchdog server has the monitored server entered in it's list of managed servers in the Enterprise Dashboard of FusionReactor. A script for the monitored server is configured that will be executed by the watchdog server if the monitored server becomes unavailable (or available again). The script is used to login to the remote computer and restart the server/instance automatically after it becomes unavailable. This approach can have the same drawbacks as mentioned in the Local Server Simple Watchdog section before.

Cross Monitoring

All servers run on different machines that can reach each other over the network. One instance/server is created for the task of Enterprise Monitor to monitor the other Operational servers. Every operational server is added to the Enterprise Monitor's list of managed servers in the Enterprise Dashboard of FusionReactor. For every server a script is configured that will be executed when the server becomes unavailable (or available again). The script is used to login to the remote computer and restart the operational server automatically if it becomes unavailable. One of the operational servers is also configured to monitor the Enterprise Monitor, and a script is configured to login to the Enterprise Monitor computer and restart the Enterprise Monitor instance should it become unavailable. If this way the Enterprise Monitor watches over all of the operational servers, and one of the operational Servers watches other the Enterprise Monitor. Alternatively a second Enterprise monitor could be configured with the task of watching over the first Enterprise monitor if you do not want operational servers to perform monitoring tasks.

Summary

FusionReactor Crash Protection Scripts are powerful functionality that can lead to an increase of server availability in local as well as distributed server environments. **You must however be aware of that a script that is executed by a Crash Protection rule is executed by the server instance that is monitoring the**

problem server and not by the problem server itself. If this is not the same computer, remote scripting has to be used to react to the situation in an appropriate way. Also note, that you should not use scripting in a single instance (self-monitoring) scenario, because the script may not reliably start.

For more information on this topic, please see "Using FusionReactor Enterprise Scripting".

Metrics

System Metrics

Longest Requests

Slow Requests

Longest JDBC

Metrics Settings

System Metrics

If Metrics are enabled (see Metrics Settings) then the Request Metrics page will display various statistics about this server. These metrics are split into several sections:

Name	Description
Graphs	Here you can see 6 graphs: Request Load, Request Time, JDBC Load, JDBC Request Time, Memory Demand, and CPU. Click on any graph to maximize it and click on a maximized graph to return to the standard view. You can also flip one or all of the graphs between the minute and hour view by selecting one of the options from the little clock menu at the top right corner of each graph.You can click on the garbage can on the Memory Demand graph to request the system garbage collector.
Overview	This section shows general statistics about this server.
Recent	This section shows statistics about actions which have occurred on this server within the last 60 seconds. The Recent Slow Pages link will show you all requests from the current request history which ran longer than the slow page threshold (defined on the Metrics Settings page).
Crash Protection (Page Aborts)	This section shows how many pages have been aborted due to each of the three types of Crash Protection. Clicking on the various Crash Protection types will show you all requests from the current request history which were aborted for that reason.

Return Code	This section lists all HTTP return codes which have been returned by this server. Each return code is also a link which will show you all requests from the current request history which have that return code.

> **NOTE**
>
> Many of the values shown on this page are counts based on all pages which have been run on this server. If you see that the number of requests which returned "200 OK" is in the thousands, the number of requests actually listed will depend upon the Request History size (which is defined in Request Settings.) It is quite possible that, if enough requests have been processed since the last "500 Internal Server Error", then clicking on that link may not show any requests at all.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Also See

Metrics Settings - Enable Metrics

Longest Requests

The Longest Requests page shows a list of the longest running requests on this server. To change the size of the Longest Request list, change the History Size value on the Metrics Settings page. The Longest Requests table has the following columns:

Name	Description
Finished	The time at which the request was completed.
Ret.Code/IP	The final HTTP return code of the request and the IP address of the caller. If the request was killed or if it triggered Crash Protection then you will see a note here letting you know what happened.
ID	The request ID and the name of the thread which is servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request took to complete. If available then you will also see the amount of actual CPU time that this request required below it.

Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown: Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit).

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles. Clicking on a title for a second time will reverse the order on that column.

Managing Requests

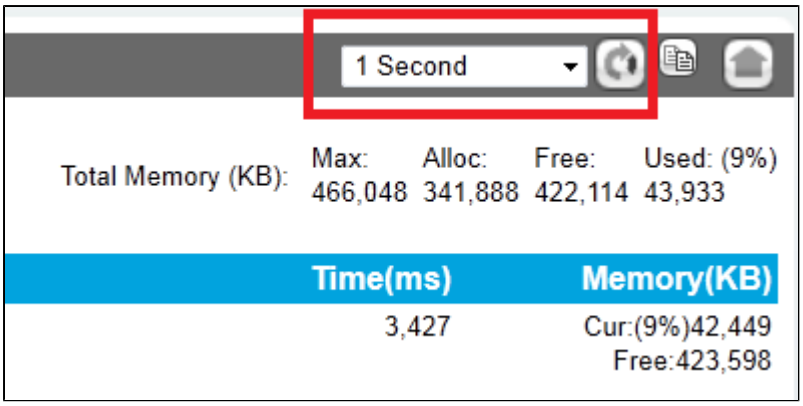
Each request listed on this page has two icons next to it. They are:

Button	Name	Description
	Request Detail	If you click this button then you will be taken to the Request Details page which gives you access to all information associated with a request. eg. Cookies, JDBC, Headers, User Trace Markers and AMF data.
	Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.

> **NOTE**
>
> Over extended periods of time you may find that this list gets filled with very long running requests. In order to see some of the medium running requests again you can empty the list by clicking on the **Clear Longest Requests** button.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.



Also See

Metrics Settings - Enable Metrics

Slow Requests

The Slow Requests page shows a list of currently running and completed requests which have taken longer to run than the current Slow Request Threshold (which is defined on the Metrics Settings page.) To change the size of the Slow Request list, change the History Size value on the Metrics Settings page. The Slow Requests table has the following columns:

Name	Description
Started *or* Finished	Currently running requests show the time at which the request was accepted. Finished requests show the time at which the request completed.
IP *or* Ret.Code/IP	The final HTTP return code of the request is shown for completed requests. Both current and completed requests show the IP address of the caller. If a completed request was killed or if it triggered Crash Protection then you will see a note here letting you know what happened.
ID	The request ID and the name of the thread which is/was servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request has taken/took to complete. If available then you will also see the amount of actual CPU time that this request used below this.

Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown:Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit).

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles. Clicking on a title for a second time will reverse the order on that column.

Managing Requests

Each currently running request listed on this page has four icons next to it. They are:

Button	Name	Description
	Stack Trace	Clicking this button will take you to the Stack Trace page, allowing you to see the current Java stack trace of the request.The request must be being currently executed to be presented with this option.
	Request Detail	Clicking this button will take you to the Request Details page, allowing you to access all information associated with a request; e.g cookies, JDBC, headers, User Trace Markers and AMF data.

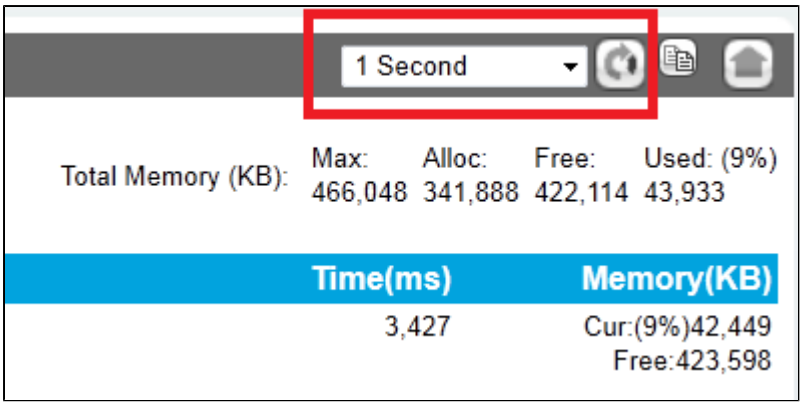
	Kill Request	Clicking this button allows you to manually kill any tracked request. You will be asked if you are sure. Click on OK to kill the request. **WARNING** Although FusionReactor first attempts to signal the request to let it exit safely, if this times out or you skip the timeout, then the request will be killed in a very abrupt manner. Please be aware that in some very limited cases, this can cause the JVM to become unstable.
	Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.
	Stop Tracking	Clicking this button will cause FusionReactor to stop tracking this request for certain metrics (slowest requests, request runtime. request activity). Once this button has been pressed, it is grayed and disabled. There is no way to reverse this operation.

Clear

The Slow Requests list can be cleared completely by clicking the **Clear History** button.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.



Also See

Metrics Settings - Enable Metrics

Longest JDBC

The Longest JDBC page shows a list of the longest running JDBC requests on this server. To change the size of the Longest Request list, change the History Size value on the Metrics Settings page. The Longest JDBC table has the following columns:

Name	Description
SQL	This is the SQL which actually ran followed by a JDBC stack trace (the list of CFCs and tags which were called to get to this query) and finally you will see the URL which was called from the browser. If the stack trace is longer than one line then it will be displayed collapsed. Click on the line or the arrow to the left of the line to expand the stack trace.
Data Source	If you are using a wrapped data source and have named it (Refer to the JDBC Driver Wrapper User Guide) then that name will appear here.
Start Time	This is the time at which this JDBC query began.
Total Time (ms)	This is the amount of time it took CF to complete the query.
DB Time (ms)	This is the amount of time the database actually spent processing the query.
Row Count	This is the number of rows which was returned by the query.

Managing Requests

Each JDBC request listed on this page has an icon next to it:

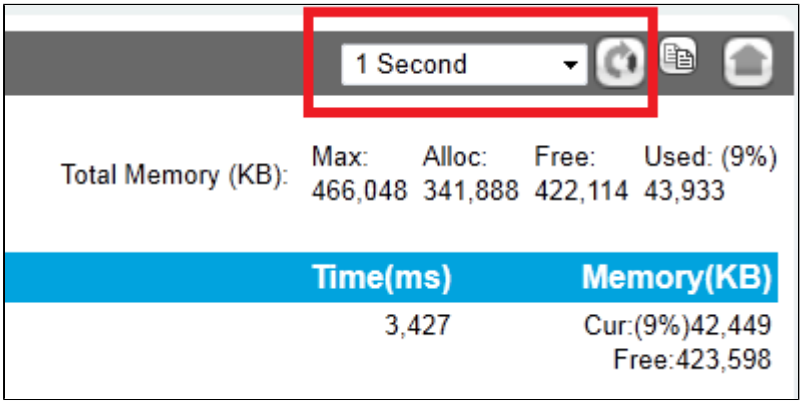
Button	Name	Description

	Request Detail	If you click this button then you will be taken to the Request Details page for the page which ran this JDBC request. This will give you access to all information associated with that request. eg. Cookies, JDBC, Headers, User Trace Markers and AMF data.

> **NOTE**
>
> Over extended periods of time you may find that this list gets filled with very long running JDBC requests. In order to see some of the medium running requests again you can empty the list by clicking on the **Clear Longest JDBC** button.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.



Also See

Metrics Settings - Enable Metrics

Metrics Settings

The Metrics Settings page lets you modify various elements relating to Metrics. The page has the following items:

Metrics:

Configuration	Description
Enable Metrics	Enable Metrics in order to access the System Metrics , Longest Requests,Slow Requests and Longest JDBC pages. You can reset all metrics by disabling and re-enabling Metrics.

Slow and Longest Running Requests:

Configuration	Description
History Size	This is the number of requests which will be shown on the Longest Requests and Slow Requests pages.
Slow Request Threshold (s)	If a request takes longer than this number of seconds then it is considered a slow running request. It will be displayed on the Slow Requests page and the slow page count on the System Metrics page will be increased. **NOTE** *This value does NOT effect that used in Crash Protection to decide Runtime Timeout.*

Slow and Longest Running Requests:

Configuration	Description
History Size	This is the number of JDBC requests which will be shown on the Longest JDBC page.

Resources

Request Activity Graph

Request Time Graph

JDBC Activity Graph

JDBC Time Graph

Memory Graphs

- Memory Heap
- Heap and Non Heap
- Types of Memory Graphs

CPU Graph

List All Threads

Resource Settings

Resource Log (resource.log)

Memory Logs

- Memory Heap and Non Heap Log (memorysummary.log)
- Types of Memory Logs

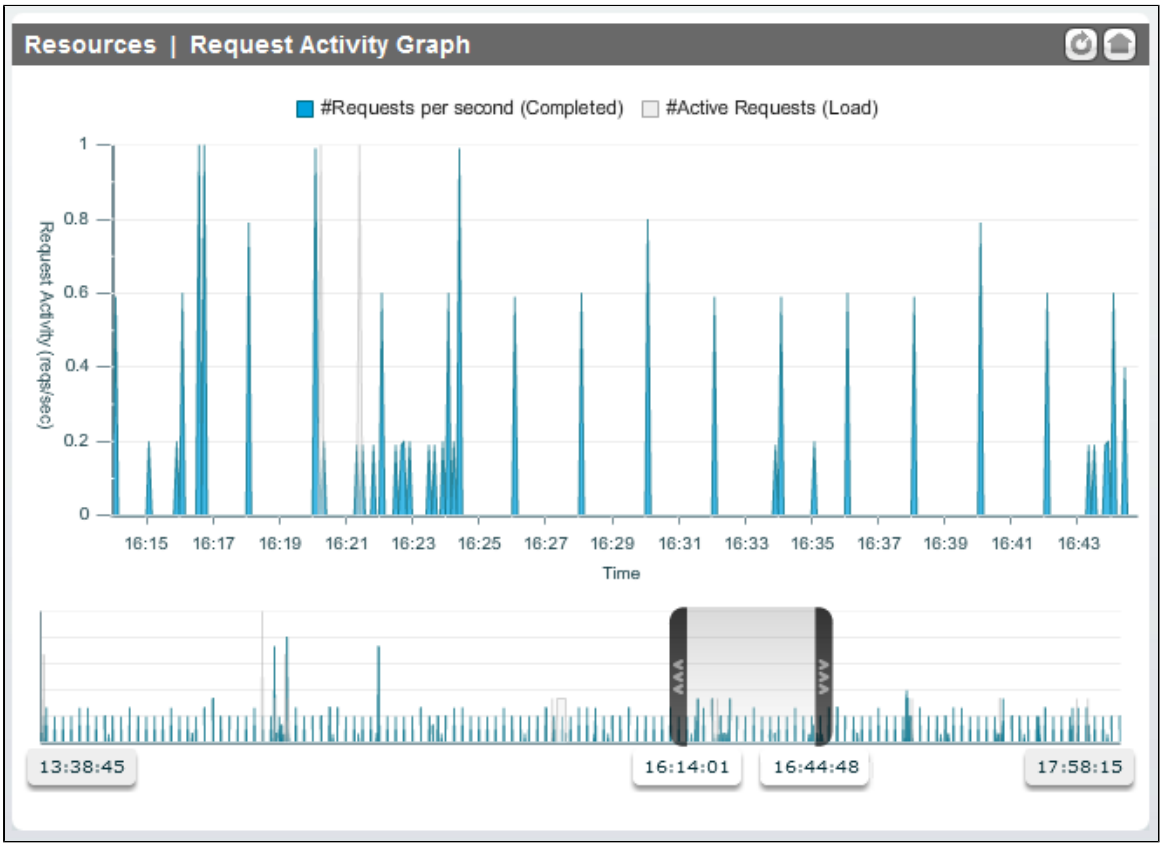
CPU Thread Log (cputhread.log)

Request Activity Graph

The Request Activity Graph is a graph of the request activity samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the number of requests completed per second since the previous sample on the graph (Blue) and the number of requests running at the time that the sample is taken (Gray). This graph gives you insight into the server Request activity; if the number of active requests stays high and constant then your server is under load. It could be that the requests are taking too long to execute for the load being placed upon the server.

Placing your mouse pointer on a data point will show a tool-tip with details about that sample.



Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.



NOTE

After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

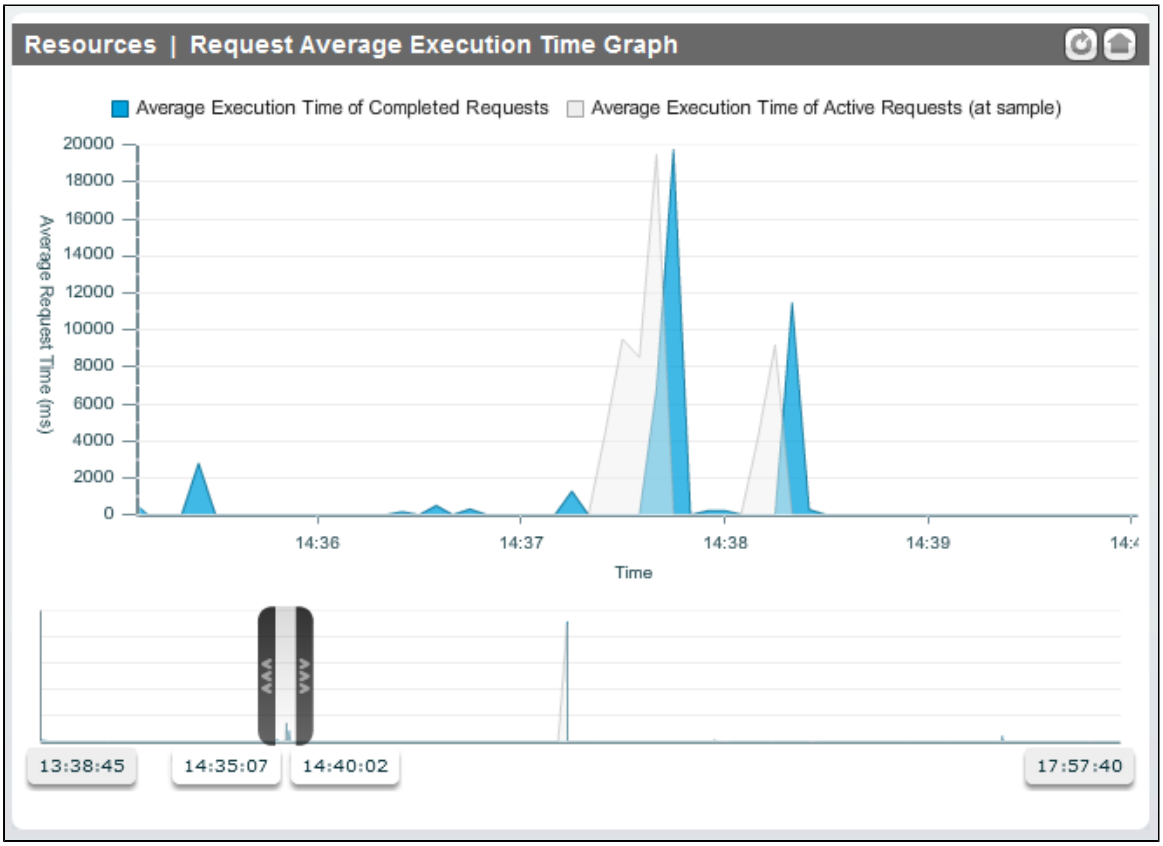
Importing and Graphing Data in Excel

Request Time Graph

The Request Average Execution Time Graph is a graph of the request performance samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the average execution time of the completed requests since the previous sample on the graph (Blue) and the average execution time of active requests at the point that the sample is taken (Gray). This graph gives you insight into the server Request performance; if the average execution time of the active requests is continuing to increase it could be that some requests are hanging.

Placing your mouse pointer on a data point will show a tooltip with details about that sample.



Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.



> **NOTE**
>
> After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

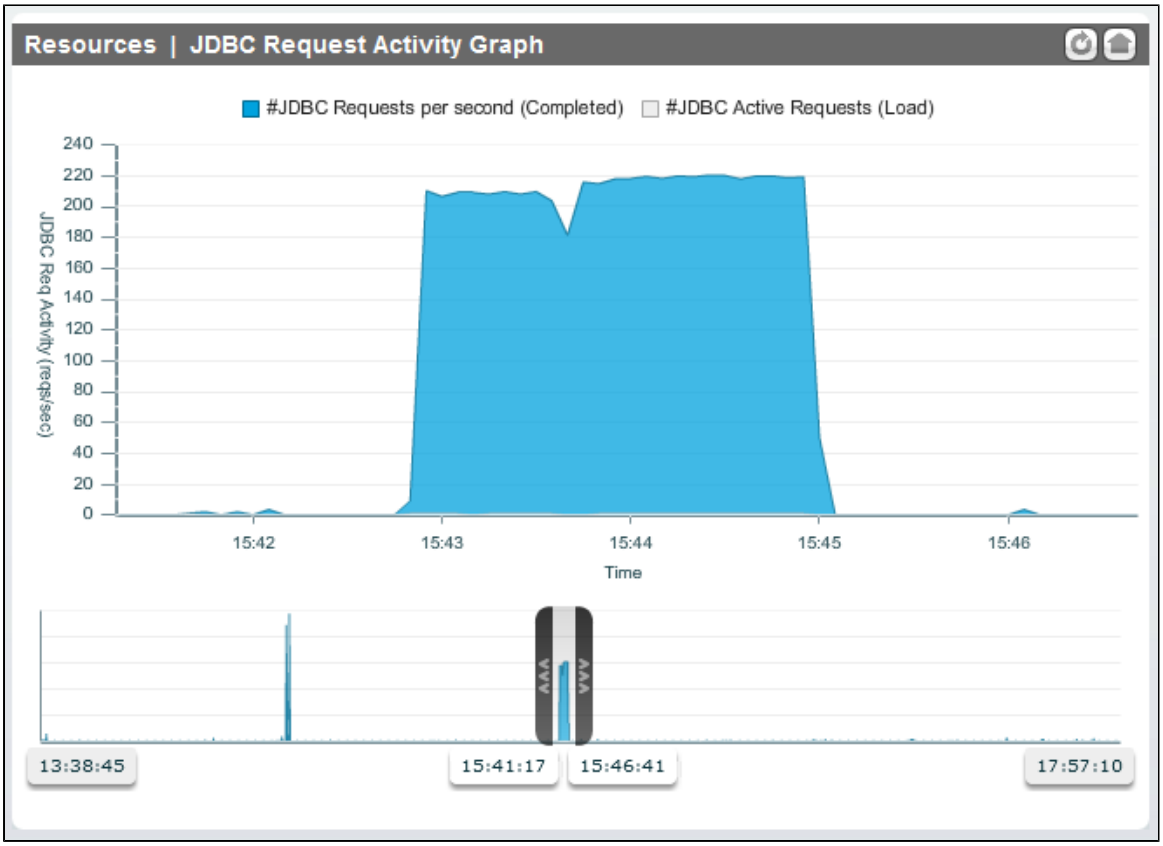
Importing and Graphing Data in Excel

JDBC Activity Graph

The JDBC Request Activity Graph is a graph of all the JDBC request activity samples currently stored by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the number of JDBC requests completed per second since the previous sample on the graph in blue, and the number of requests running at the time that the sample is taken (Gray). This graphs give you insight into the JDBC Request activity; if the number of active requests stays high and constant then your server is under load. It could be that the requests are taking too long to execute for the load being placed upon the server.

Placing your mouse pointer on a data point will show a tooltip with details about that sample.



Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.



> **NOTE**
>
> After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

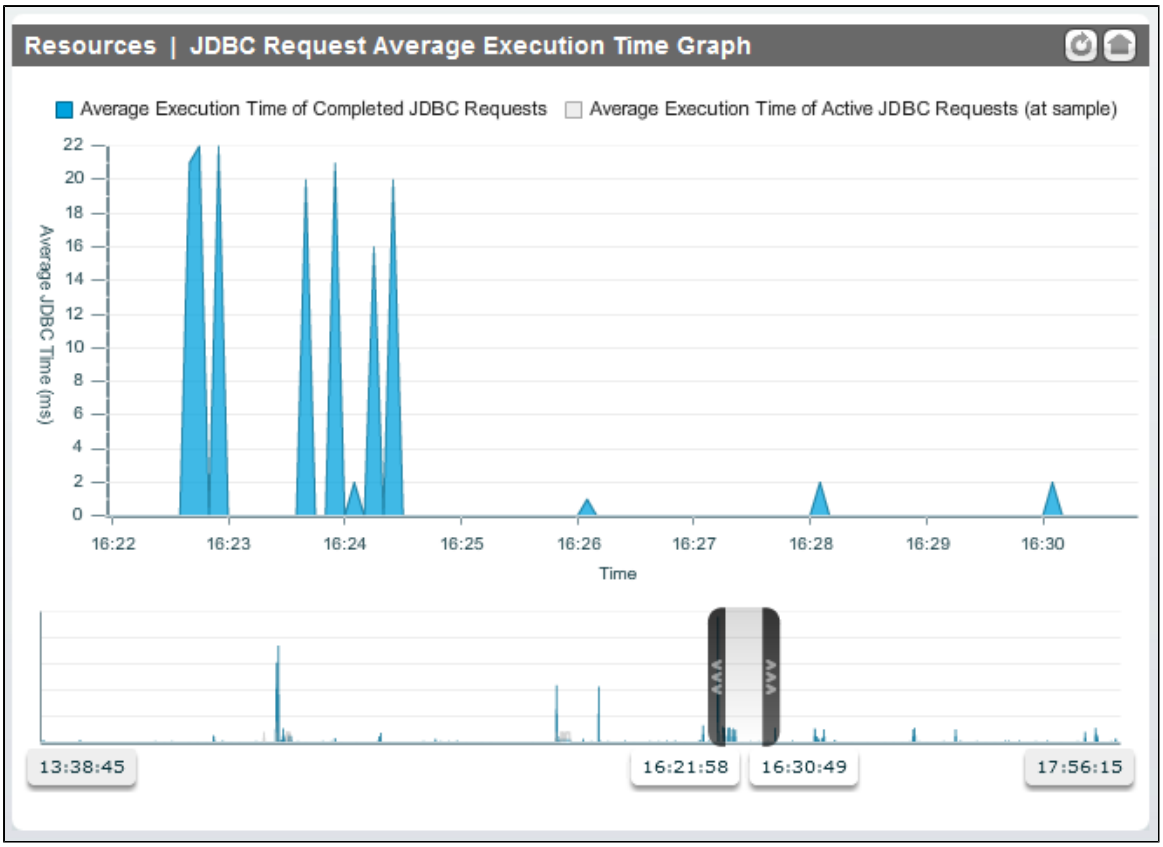
Importing and Graphing Data in Excel

JDBC Time Graph

The JDBC Request Average Execution Time Graph is a graph of all the JDBC request performance samples currently stored by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the average execution time of the completed JDBC requests since the previous sample on the graph (Blue) and the average execution time of active JDBC requests at the point that the sample is taken (Gray). This graph gives you insight into the JDBC Request performance; if the average execution time of the active JDBC requests is continuing to increase it could be that some JDBC requests are hanging.

Placing your mouse pointer on a data point will show a tooltip with details about that sample.



Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.



NOTE

After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

Importing and Graphing Data in Excel

Memory Graphs

Memory Heap

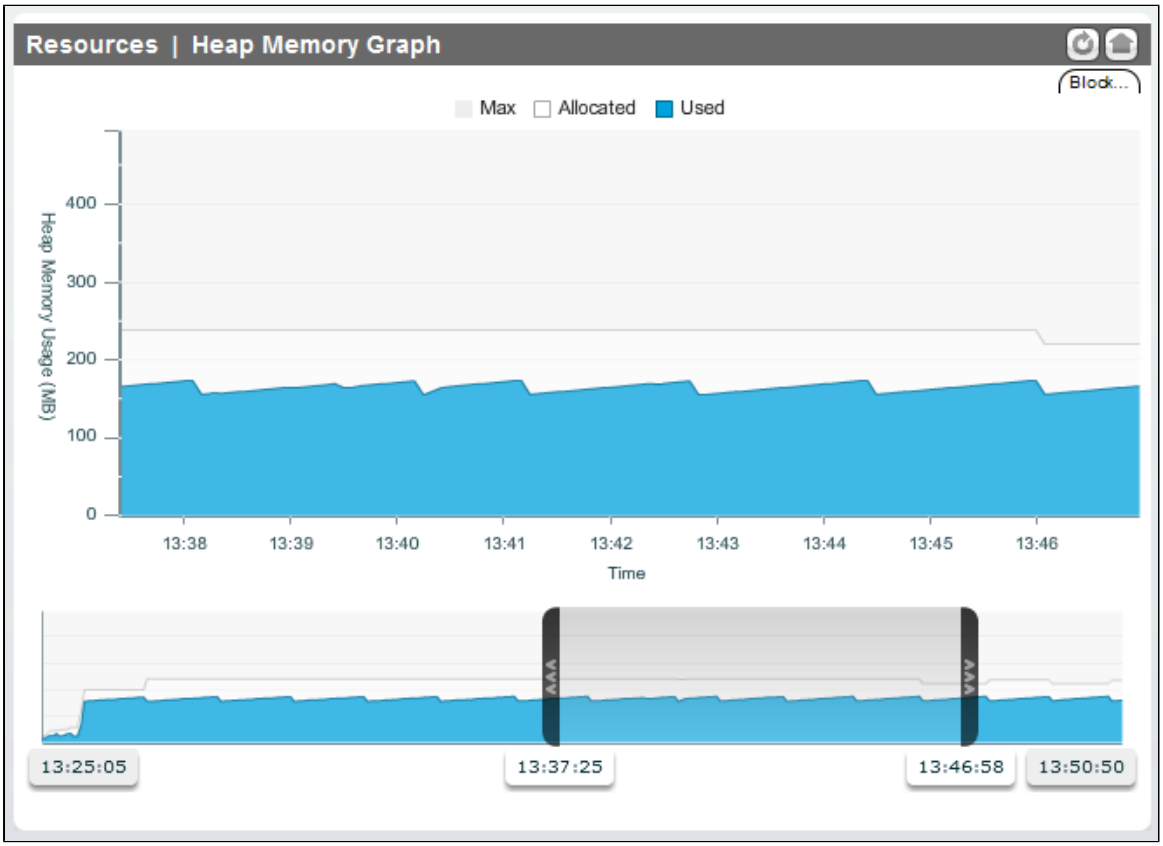
Heap and Non Heap

Types of Memory Graphs

Memory Heap

The Memory Heap Graph will show you a graph of the memory samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

From this graph you can see the amount of unallocated memory, the amount of allocated memory and the amount of memory actually being used when the sample was taken.



Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.



Why would I want to see a memory heap?

Because of its visual nature, the memory heap graph can very quickly give you a good idea how the system is performing. Not only that but, because it can cover larger periods of time, it can warn you of trends which could become a problem in the future. Trends are much more difficult to spot if you are only checking the system periodically and are only looking at the health of the system right at this point in time.

> **NOTE**
>
> After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

Importing and Graphing Data in Excel
Resource Settings

Heap and Non Heap

The Heap and Non Heap Graph will show you a graph of the memory samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

From this graph you can see the amount of unallocated memory, the amount of allocated memory and the amount of memory actually being used when the sample was taken.

Heap memory

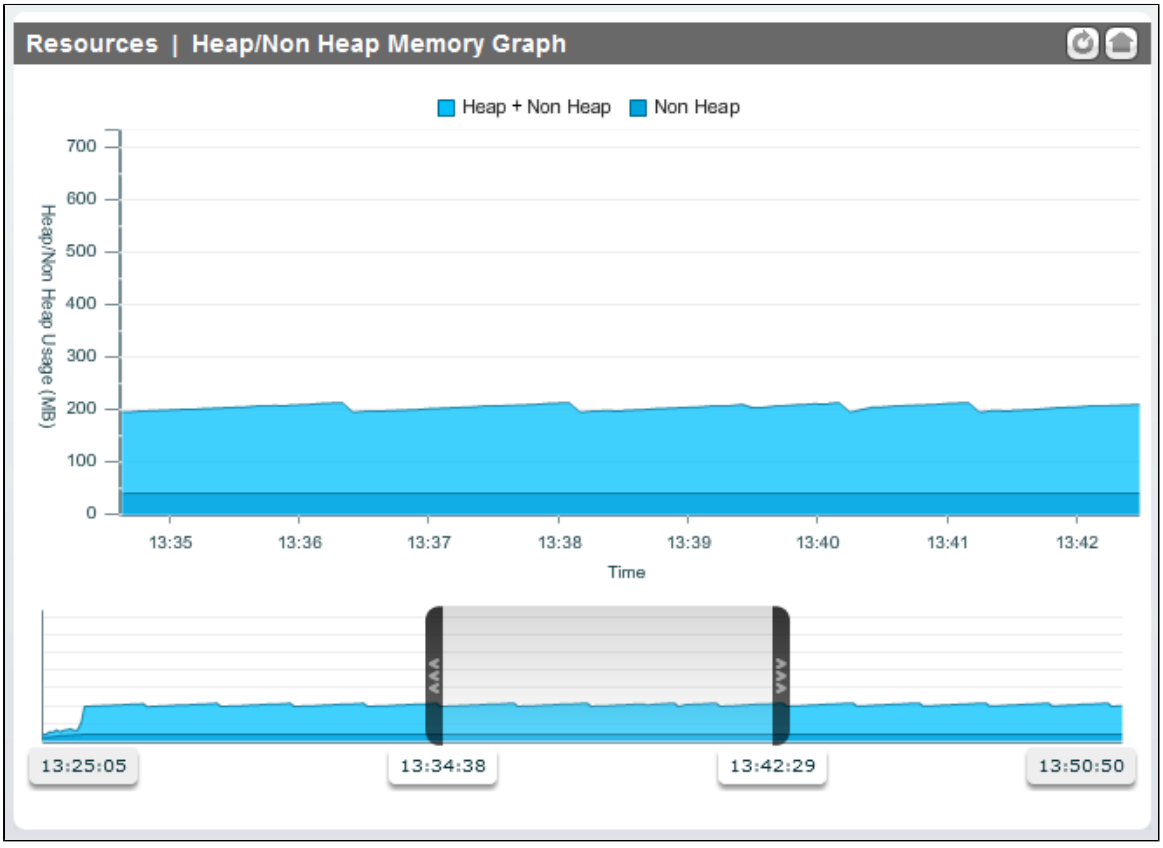
The heap memory is the runtime data area from which the Java VM allocates memory for all class instances and arrays. The heap may be of a fixed or variable size. The garbage collector is an automatic memory management system that reclaims heap memory for objects.

- Eden Space: The pool from which memory is initially allocated for most objects.
- Survivor Space: The pool containing objects that have survived the garbage collection of the Eden space.
- Tenured Generation: The pool containing objects that have existed for some time in the survivor space.

Non-heap memory

Non-heap memory includes a method area shared among all threads and memory required for the internal processing or optimization for the Java VM. It stores per-class structures such as a runtime constant pool, field and method data, and the code for methods and constructors. The method area is logically part of the heap but, depending on the implementation, a Java VM may not garbage collect or compact it. Like the heap memory, the method area may be of a fixed or variable size.

- Permanent Generation: The pool containing all the reflective data of the virtual machine itself, such as class and method objects. With Java VMs that use class data sharing, this generation is divided into read-only and read-write areas.
- Code Cache: The HotSpot Java VM also includes a code cache, containing memory that is used for compilation and storage of native code.



Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.



Why would I want to see a Heap and Non Heap Memory?

This graph is very efficient as it allows you to quickly analyze non heap memory, heap memory and a combination of both. Not only that but, because it can cover larger periods of time, you can see if anything is becoming critical meaning you can act before it is to late. This kind of trend is much more difficult to spot if you are only checking the system periodically and are only looking at the health of the system right at this point in time.

> **NOTE**
>
> After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

Importing and Graphing Data in Excel
Resource Settings

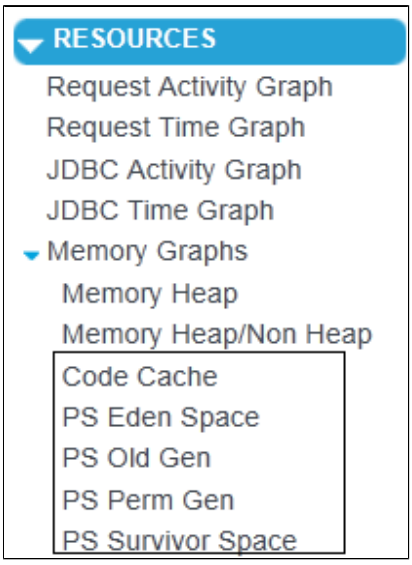
Types of Memory Graphs

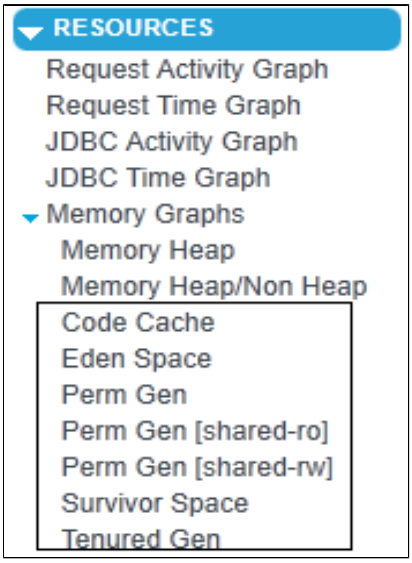
Each server (ColdFusion, Tomcat, Railo etc) has their own naming conventions for areas in memory. When an instance of FusionReactor is installed on a server you will be able to see detailed graphs for certain areas of memory that the server has. The naming conventions used by your server for areas of memory will be listed in FusionReactor meaning you won't have to translate them to terms that relate to your server. For example an area in memory in ColdFusion 9 is PS Eden Space, so there is a graph and a log for PS Eden Space.

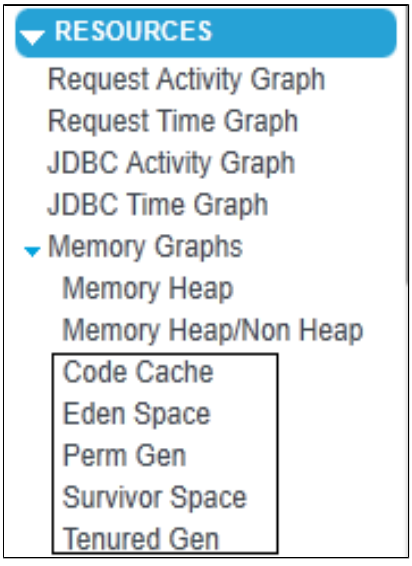
> **NOTE**
>
> For FusionReactor to be able to detect some of the memory spaces this will only be available on JMX architectures (you must have Java 1.5 or above installed).

This page lists some of the naming conventions used be a couple of servers to give you an understanding of how it works.

Server (s)	Log	Image

ColdFusion 10 ColdFusion 9 ColdFusion 8	• Code Cache (non-heap memory) • PS Eden Space (heap memory) • PS Old Gen (heap memory) • PS Perm Gen (non-heap memory) • PS Survivor Space (heap memory)	

Server (s)	Log	Image
Tomcat 7	• Code Cache • Eden Space • Perm Gen • Perm Gen [shared-ro] • Perm Gen[shared-rw] • Survivor Space • Tenured Gen	

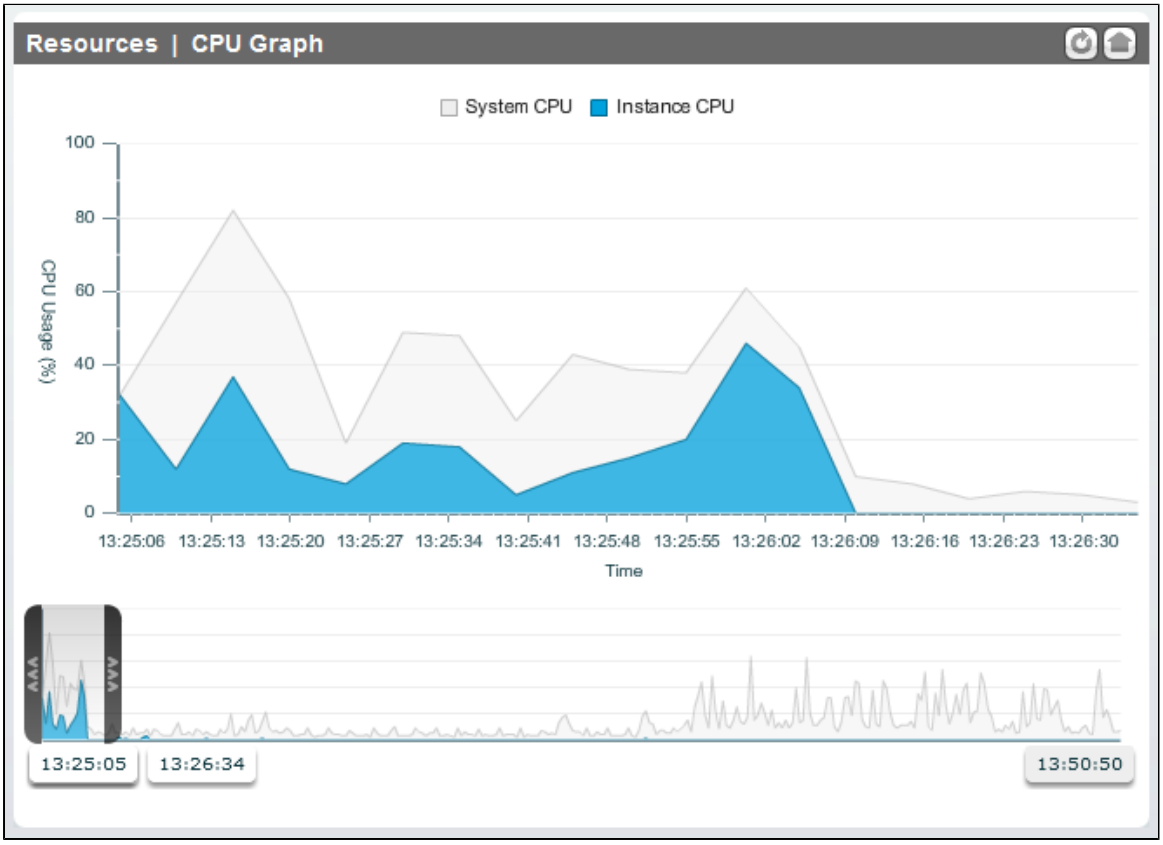
Server (s)	Log	Image
Server Names: Railo 3	• Code Cache • Eden Space • Perm Gen • Survivor Space • Tenured Gen	

Will I still get detailed information in my logs and graphs?

As you can see every server has different names for areas of memory so there is no hard rule about what graphs will be displayed in FusionReactor. There will be graphs and logs available for areas of memory in your server. **So the short answer to the question above is yes.** For further information on memory logs, see the Types of Memory Logs section.

CPU Graph

The CPU Graph will show you a graph of the CPU samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds). The System CPU is presented in Gray and the Instance CPU is presented in blue.



Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.



> **NOTE**
>
> After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

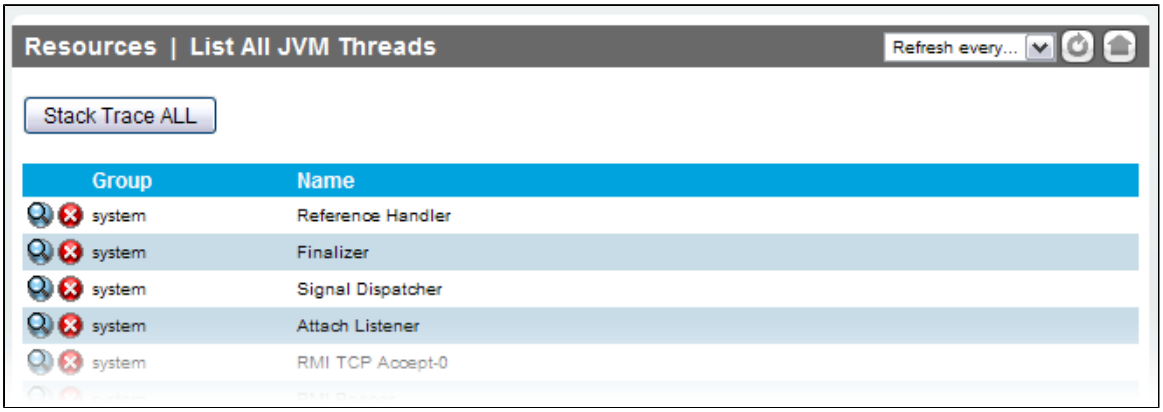
Also See

Importing and Graphing Data in Excel

List All Threads

The List All Threads page shows all threads running within the JVM. The threads table has the following columns:

Name	Description
Group	This is how Threads are organized within the JVM.
Name	This is the name of a Thread.



Managing Threads

Each thread listed on this page has two icons next to it. They are:

Button	Name	Description
	Stack Trace	If you click this button then you will be taken to the Stack Trace page which lets you see exactly what is currently going on inside a thread.
	Kill Thread	Clicking on this link will prompt you if you are sure. OK the dialog to kill the thread.

WARNING

All JVM threads are listed in this page, including those pertaining to JVM system and J2EE container functions. Killing a thread is a last resort to solving problems and we must caution you that, while FusionReactor does allows you to stop any thread, it is highly inadvisable to stop any threads in the "system" thread group, or any threads integral to the running of your J2EE container.

Refreshing the Page

At the top right of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.



Resource Settings

The Resource Settings page lets you modify various elements relating to resource sampling and logging. The page has the following items:

Resource Sampling (CPU/Memory):

Configuration	Description
Initial Delay (ms)	This value defines how long after startup FusionReactor will wait before starting to periodically gather resource information. This is done so that the server can recover from the extra overhead of startup and your graphs will represent normal operating behavior.
Sample Interval (ms)	Change this value to alter the amount of time between successive resource samples.
History Size	This value defines the maximum number of resource samples stored in memory. Once the history size reaches its maximum, old values will be removed to make room for new ones.

Java 1.5+ Logging (CPU-Thread Status/Memory Spaces)

Configuration	Description

CPU/Thread Status	Setting this to Enabled will allow CPU/Thread information to be logged, setting it to disabled will not log the information. **WARNING** Enabling this option will create very detailed log files, please ensure sufficient disk space is available before doing so.
Memory Space	When Enabled detailed sizing information about memory spaces will be logged.

NOTE

To Enable the two options above you must be running Java 1.5 or above. See the Sun Website for information on how to download the latest version of Java.

Resource Logging (CPU/Memory/Request/JDBC):

Configuration	Description
Resource Logging	If Resource Logging is enabled then resource information will be written to the Resource Log.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the Resource Log from inside FusionReactor.

Resource Log (resource.log)

The Resource Logs (a set of rotating log files accessible over the Resource Settings page) are simple space-separated files so they can be easily imported into database or spreadsheet applications. The file is called **resource.log**, the bottom part of the page tells you where you can locate this file on your hard drive.

The log file is formatted as follows:

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this resource probe was triggered.

Time (Formatted)	2 [B]	The time at which this resource probe was triggered.
Date/Time (ms)	3 [C]	This is the millisecond time at which at which this resource probe was triggered.
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
MEM	6 [F]	Fixed Separator
Used Memory (percentage)	7 [G]	The amount of used memory (expressed as a percentage.).
Max Memory (KB)	8 [H]	The total amount of physical memory available to this instance.
Used Memory (KB)	9 [I]	The amount of used memory.
Total Memory (KB)	10 [J]	The amount of allocated memory.
Free Memory (KB)	11 [K]	The amount of free memory (within the allocated block).
CPU	12 [L]	Fixed Separator
CPU Time (percentage)	13 [M]	This is the amount of CPU time used by CF since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.
Total CPU (percentage)	14 [N]	This is the amount of CPU time used as a whole (all running applications) by the server since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.
REQ	15 [O]	Fixed Separator
Total Request Count	16 [P]	The total number of requests that have been called since the last restart.

Total Finished Request Count	17 [Q]	The total number of requests that have been completed since the last restart.
Request Count	18 [R]	The number of requests that were running at the time of this log message.
Queued Request Count	19 [S]	The number of requests that were queued at the time of this log message.
Finished Request Count	20 [T]	The number of requests that were completed since the last log message.
Requests per second	21 [U]	The number of requests per second that where being served at the time of this log message.
Average Request time (ms)	22 [V]	The average request execution time at the time of this log message.
JDBC	23 [W]	Fixed Separator
Total JDBC Count	24 [X]	The total number of JDBC requests that have been called since the last restart.
Total Finished JDBC Count	25 [Y]	The total number of JDBC requests that have been completed since the last restart.
JDBC Count	26 [Z]	The number of JDBC requests that were running at the time of this log message.
JDBC Finished	27 [AA]	The number of JDBC requests that were completed since the last log message.
JDBC Requests per second	28 [AB]	The number of JDBC requests per second that where being served at the time of this log message.
Average JDBC time	29 [AC]	The average JDBC request execution time at the time of this log message.

Deleting the Resource Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the

Resource Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See

Log Settings
Importing and Graphing Data in Excel

Memory Logs

The Memory Logs section contains two pages. One explains the memorysummary.log which is a basically a log containing heap and non heap memory information. The other page talk about how every server references memory in different ways and how FusionReactor handles this and manages to create a logs that match the types of memory used by that server, see the Types of Memory Logs page for more information.

> **NOTE**
>
> Logs can be located one of two places depending on what configuration you choose in Log Settings. Both places are as follows:
>
> 1. If you choose Use Centralized Archive and Rotation Use Centralized Archive and Rotation it is located C:\FusionReactor\instance\FRAM\log[*archive_name*].
> 2. If you choose Use Traditional Per-Log Rotation it is located C:\FusionReactor\instance\FRAM\log. Notice these logs aren't archived.

Memory Heap and Non Heap Log (memorysummary.log)

Types of Memory Logs

Memory Heap and Non Heap Log (memorysummary.log)

For information on changing log settings see the Log Settings page. This log displays information about the memory heap and non heap memory. Each log entry occupies one line and contains various pieces of useful information, separated with a white space character. Below is a table explaining each element in the log. Elements are separated with a white space character. The log is called memorysummary.log.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this resource probe was triggered.
Time (Formatted)	2 [B]	The time at which this resource probe was triggered.
Date/Time (ms)	3 [C]	This is the millisecond time at which at which this resource probe was triggered.
Version	4 [D]	Version number of this log entry
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.

Heap Memory Allowance (Percentage)	6 [F]	The amount of heap memory that can be used for memory management.
Heap Memory Used (Bytes)	7 [G]	The amount of heap memory used.
Heap Memory Committed Bytes)	8 [H]	The amount of memory that can be used for the heap by the Java Virtual Machine.
Heap Memory Requested (Bytes)	9 [I]	The amount of heap memory that the Java Virtual Machine initially requests from the operating system.
Heap Memory Allowance (Bytes)	10 [J]	The amount of heap memory that can be used for the heap by the Java Virtual Machine.
Non Heap Memory Allowance (Percentage)	11 [K]	The amount of non heap memory that can be used for memory management.
Non Heap Memory Used (Bytes)	12 [L]	The amount of non heap memory used.
Non Heap Memory Committed (Bytes)	13 [M]	The amount of non heap memory that can be used for the heap by the Java Virtual Machine.
Non Heap Memory Requested (Bytes)	14 [N]	The amount of non heap memory that the Java Virtual Machine initially requests from the operating system.
Non Heap Memory Allowance (Bytes)	15 [O]	The amount of non heap memory that can be used for memory management.
Object Finalization Count	16 [P]	The number of objects waiting to have their finalize method called. This is the method that is called before the object is destroyed.

Deleting the Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the Memory Heap and Non Heap Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Broken Down Memory Logs

The heap memory and non heap memory can be further broken down to specific areas. Every server handles areas of memory differently and refers to them using terms other servers may not use. FusionReactor recognizes this and creates logs for all the different areas of memory that particular server supports, furthermore the logs are named after the area of memory that the server uses making it easier for you to locate your logs. For more information on broken down memory logs see the Types of Memory Logs section.

Also See

Log Settings
Heap and Non Heap Graphs

Types of Memory Logs

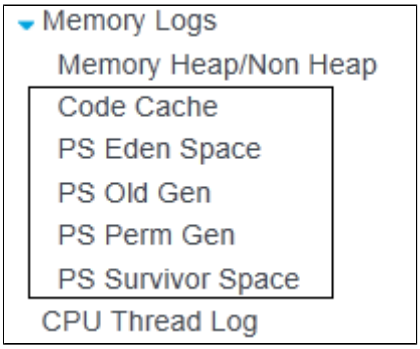
> **NOTE**
>
> For FusionReactor to be able to detect what areas of memory are referred to in your server you must have Java 1.5 or above installed.

> **NOTE**
>
> Each server (ColdFusion, Tomcat, Railo etc) has their own naming conventions for areas in memory. When an instance of FusionReactor is installed on a server you will be able to see detailed logs for certain areas of memory that the server has. The naming conventions used by your server for areas of memory will be listed in FusionReactor meaning you won't have to translate them to terms that relate to your server. For example an area in memory in ColdFusion 9 is PS Old Gen, so there is a log and a graph for PS Old Gen.

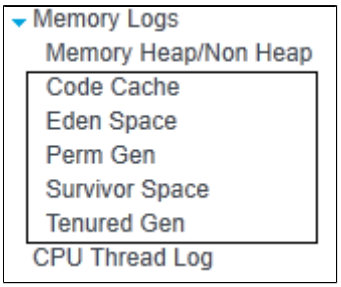
This page lists some of the naming conventions used by a couple of servers to give you an understanding of how it works.

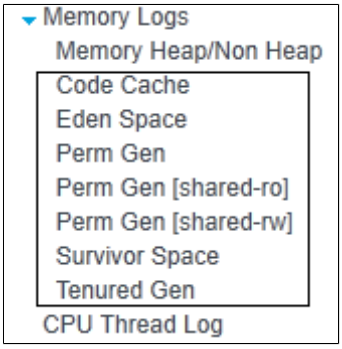
Log File Names

The following tables contain the name of the server along with the log files and a screen shot showing how to view the log files in FusionReactor.

Server (s)	Log Files	Image
ColdFusion 10 ColdFusion 9 ColdFusion 8	• memory-Code-Cache.log • memory-PS-Eden-Space.log • memory-PS-Old-Gen.log • memory-PS-Perm-Gen.log • memory-PS-Survivor-Space.log	

Server (s)	Log Files	Image

Server Names: Railo 3	• memory-Code-Cache.log • memory-Eden-Space.log • memory-Perm-Gen.log • memory-Survivor-Space.log • memory-Tenured-Gen.log	

Server (s)	Log Files	Image
Server Names: Railo 7	• memory-Code-Cache.log • memory-Eden-Space.log • memory-Perm-Gen.log • memory-Perm-Gen-[shared-ro].log • memory-Perm-Gen-[shared-rw].log • memory-Survivor-Space.log • memory-Tenured-Gen.log	

Log fields

All areas of memory are stored in logs in a similar fashion to each other with the only difference being an attribute that identifies the log type. Below is the standard layout for any kind of memory log file (excluding Memory Heap / Non Heap log).

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this resource probe was triggered.
Time (Formatted)	2 [B]	The time at which this resource probe was triggered.
Date/Time (ms)	3 [C]	This is the millisecond time at which at which this resource probe was triggered.
Version	4 [D]	Version number of this log entry
Server Startup Time (ms)	5 [E]	Start up time of the application server in milliseconds (ms) since epoch.
Name	6 [F]	The name that represents the memory pool.
Log File Type	7 [G]	The types of memory that the log relates to (heap or non heap).

Memory Collection Used (Percent)	8 [H]	The amount of collection memory used.
Memory Collection Usage Used (Bytes)	9 [I]	The amount of collection memory used.
Memory Collection Usage Committed (Bytes)	10 [J]	The amount of collection memory committed for the JVM to used.
Memory Collection Usage Requested (Bytes)	11 [K]	The amount of collection memory initially requested from the OS by the JVM.
Memory Collection Usage Allowance (Bytes)	12 [L]	The total amount of collection memory that can be used for memory management.
Collection Usage Threshold Count	13 [M]	The number of times that the JVM has detected that the memory usage has reached or exceeded the threshold.
Collection Usage Threshold (bytes)	14 [N]	The collection usage threshold value of the memory pool.
Is Collection Threshold Exceeded	15 [O]	If the memory usage of the memory pool has reached or exceeded the threshold count. If it has exceeded the value will be 1, otherwise it will be 0.
Memory Usage Used(Percent)	16 [P]	The amount of memory used.
Memory Usage Used (Bytes)	17 [Q]	The amount of memory used.
Memory Usage Committed (Bytes)	18 [R]	The amount of memory committed for the JVM to used.
Memory Usage Requested (Bytes)	19 [S]	The amount of memory initially requested from the OS by the JVM.
Memory Usage Allowance (Bytes)	20 [T]	The total amount of memory that can be used for memory management.
Usage Threshold Count	21 [U]	The number of times that the JVM has detected that the memory usage has reached or exceeded the threshold.
Usage Threshold (bytes)	22 [V]	The usage threshold value of the memory pool.

Is Usage Threshold Exceeded	23 [W]	If the memory usage of the memory pool has reached or exceeded its usage threshold value. If it has exceeded the value will be 1, otherwise it will be 0
Peak Memory Usage Used (Percent)	24 [X]	The peak amount of memory used.
Peak Memory Usage Used (Bytes)	25 [Y]	The peak amount of memory used.
Peak Memory Usage Committed (Bytes)	26 [Z]	The peak amount of memory committed for the JVM to use.
Peak Memory Usage Requested (Bytes)	27 [AA]	The peak amount of memory initially requested from the OS by the JVM.
Peak Memory Usage Allowance (Bytes)	28 [AB]	The peak amount of memory that has been used for memory management

What is the difference between "Usage" and "Collection Usage"

The table contains two terms that are commonly used, one being called "Usage" and the other "Collection Usage". **Usage** is basically the amount of memory that has been allocated, this can be heap memory or non heap memory. **Collection Usage** relates to the garbage collector. Whenever it runs and cleans up memory, some of that memory will still be in use and that's the memory that the table refers to when it talks about **Collecti on Usage**.

Summary

As you can see every server has different names for areas of memory so there is no hard rule about what logs will be displayed in FusionReactor. There will be logs available for areas of memory in your server.

Also See

Log Settings

CPU Thread Log (cputhread.log)

To enable CPU Thread logs enable the option on the Resource Settings page. The log allows you to view information about threads that ran on the servers CPU. Each log entry occupies one line and contains various pieces of useful information. Below is a table explaining each element in the log. Elements are separated with a white space character.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this resource probe was triggered.
Time (Formatted)	2 [B]	The time at which this resource probe was triggered.

Date/Time (ms)	3 [C]	This is the millisecond time at which at which this resource probe was triggered.
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of FusionReactor server in milliseconds (ms) since epoch.
Thread Name	6 [F]	The name of the thread.
Thread ID	7 [G]	The thread ID.
CPU Time (nanoseconds)	8 [H]	The amount of time a thread has been alive for. A thread is alive if it's running and has not died.
User Time (nanoseconds)	9 [I]	The amount of time that a thread has executed in user mode.
CPU Time Delta (nanoseconds)	10 [J]	The difference between the the current CPU Time and the previous CPU Time attributes. Current being the thread one in front of the previous in the log entries.
User Time Delta (nanoseconds)	11 [K]	The difference between the the current User Time and the previous User Time attributes. Current being the thread one in front of the previous in the log entries.
Thread Blocked Count	12 [L]	The number of times that the thread has been in the blocked state.
Thread Blocked Time (milliseconds)	13 [M]	The total amount of time the thread has been in the blocked state.
Thread Blocked Delta	14 [N]	The difference between the amount of times the current thread has been in the blocked state and the previous thread has been in the blocked state. Current being the thread one in front of the previous in the log entries.

Thread Blocked Time Delta (milliseconds)	15 [O]	The time difference between the time period the current thread has been in the blocked state and the previous thread has been in the blocked state. Current being the thread one in front of the previous in the log entries.
Lock Name	16 [P]	Monitor lock that the thread is associated with.
Lock Owner ID	17 [Q]	The ID of the thread that holds the monitor lock.
Lock Owner Name	18 [R]	The name of the thread that holds the monitor lock.
Waited Count	19 [S]	The number of times the thread has been in the waiting state.
Waited Time (milliseconds)	20 [T]	The total amount of time the thread has been in the waiting state.
Waited Count Delta	21 [U]	The difference between the amount of times the current thread has been in the waiting state and the previous thread has been in the waiting state. Current being the thread one in front of the previous in the log entries.
Waited Time Delta (milliseconds)	22 [V]	The time difference between the time period the current thread has been in the waiting state and the previous thread has been in the waiting state. Current being the thread one in front of the previous in the log entries.
Thread State	23 [W]	The thread state. States include new, runnable, blocked, waiting, timed_waiting and terminated.
Sample Period	24 [X]	The difference in time between the current resource probe being triggered and the previous resource probe being triggered. Current being the thread one in front of the previous in the log entries,

See also

Log Settings
CPU Graph
CF Realtime and Request Stats (realtimestats.log)

Crash Protection

Crash Protection Restrictions

Crash Protection Settings

Crash Protection Log (crashprotection.log)

Crash Protection Restrictions

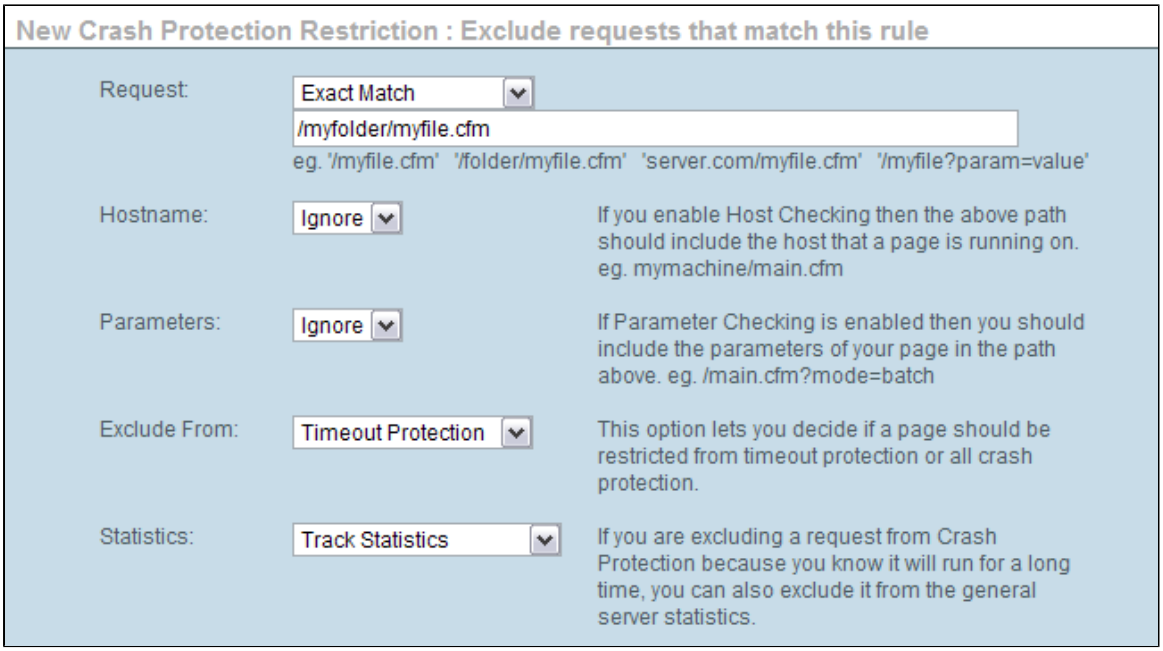
Crash Protection restrictions are designed to let you specify which pages you want to be monitored by the Crash Protection system, or alternatively, which pages you do not wish to be monitored. To select which mode the Restrictions page works in, change the Restrictions option on the Crash Protection Settings page. The Crash Protection Restrictions page has two main areas. At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for this feature.

New Crash Protection Restriction

To enter a restriction you will need to fill the following fields:

Configuration	Description
Request	Firstly select if you are creating an "Exact Match" or "Regular Expression". **Exact Match** Next enter the location of your page. By default the request will be an exact match and should look something like "/batchStats.cfm" or "/batch/dailyStats.cfm" (beginning with a '/' and without parameters.) . The text you input here will depend on the options you choose for the Hostname and Parameters fields. **Regular Expressions** A good use of Regular Expressions here would be to restrict everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Content Filters. The text you input here will depend on the options you choose for the Hostname and Parameters fields.
Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you enable Hostname checking then the request field should also include a hostname and look something like "productionMachine/batchStats.cfm". You should NOT include HTTP or HTTPS at the start of this string.

Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the Hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example "/batchStats.cfm?period=YEAR". Once again, Regular Expressions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.
Exclude From:	The meaning of this field changed depending upon what mode the Crash Protection restrictions page is working in (Crash Protection Settings page). You can either select whether the specified request will be excluded from Crash Protection, or if it should be included.
Statistics	This field is only visible when the Crash Protection restrictions page is in exclude mode (Crash Protection Settings page). If you are excluding a request from Crash Protection because you know it will run for a long time, you can also exclude it from the general server statistics.

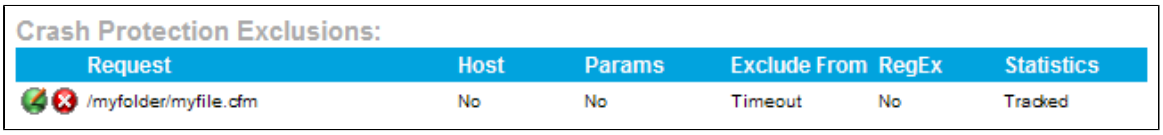


Once you submit the form the page will be refreshed and your new restriction will be listed at the bottom of the page.

Managing Existing Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

Button	Name	Description
	Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above.
	Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.



Crash Protection Settings

Unattended Monitoring via Crash Protection is a way to protect your server from various types of load. You can protect from excessive numbers of requests all being run at the same time. You can protect from individual requests running for long periods of time and you can protect from requests running when the server is low on memory. FusionReactor can notify you when any of these situations occur. The Crash Protection Settings page lets you modify various elements relating to the different forms of Crash Protection.

How do I set up Unattended Monitoring with Crash Protection?

From the FusionReactor Administrator, click "CP Settings" from the table of contents (Crash Protection section) and you will see the Crash Protection Settings screen. From here you can activate any of the crash protection triggers and set up what action should be taken if they trigger:

Crash Protection:

Configuration	Description

Timeout Protection (s)	By entering a number (in seconds) in this field you will enable Timeout Crash Protection, from that point on, any request which takes longer to execute than this time will trigger Crash Protection. To turn this feature off again, simply remove any value from the field. You can specify what should occur when Timeout Crash Protection is triggered: • **Abort and Notify** - The request will be aborted and a notification email will be sent • **Notify** - The request is left running and a warning email is sent. Requests that are aborted by crash protection can also optionally be excluded from calculation of some metrics by setting the **When Triggered** actio n: • **Continue Tracking Request** - (default) the request will continue to be tracked and incorporated in subsequent calculations • **Stop Tracking Request** - the request will *not* be incorporated in subsequent calculations. Requests which are not tracked won't be counted when calculating long-running requests, or average runtimes. Any JDBC statements run by these requests will also be ignored by tracking. The tracking policy set here will apply whether or not the request kill was successful. Such requests will be labelled as **Untracked** in the user interface, and will not cause a "red cube" condition on the **Enterp rise Dashboard** (for users with applicable licenses).
Memory Protection (%)	By entering a number (percentage) in this field you will enable Memory Crash Protection. From that point on, if the amount of free memory on the server is less than this number when a request arrives, Crash Protection will be triggered. To turn this feature off again, simply remove any value from the field. You can specify what should occur when Memory Crash Protection is triggered: • **Abort and Notify** - The request will be aborted and a notification email will be sent • **Queue and Notify** - The request will be added to a queue in order to be executed once the server has recovered and a notification email will be sent. You can specify the queue timeout in the Survival Strategy section of this page. • **Notify** - The request is left running and a warning email is sent.

Request Protection	By entering a number in this field you will enable Request Quantity Crash Protection. From that point on, if the amount of currently running requests on the server is greater than this number when a request arrives, Crash Protection will be triggered. To turn this feature off again, simply remove any value from the field. You can specify what should occur when Request Quantity Crash Protection is triggered: • **Abort and Notify** - The request will be aborted and a notification email will be sent • **Queue and Notify** - The request will be added to a queue in order to be executed once the server has recovered and a notification email will be sent. You can specify the queue timeout in the Survival Strategy section of this page. • **Notify** - The request is left running and a warning email is sent.

NOTE

In order to receive notification email you must set up your Email Server on the FusionReactor Settings page. You can also change the behavior of the Request Abort and Request Queue in the Survival Strategy part of this page.

Crash Protection Restrictions:

Configuration	Description
Restrictions	From here you can enable or disable the Crash Protection Restrictions feature. When restrictions are enabled you can define which pages will be run through the Crash Protection system and which will be left to run normally. You can define individual restrictions on the CP Restrictions page.
Behavior	This value defines how Crash Protection Restrictions behave. • **Protect Matching Requests** - No requests will be subject to protection (they will never be aborted or timed out) unless they are matched by a Restriction rule as defined on the CP Restrictions page. • **Exclude Ignore Requests** - All requests will be subject to protection (they may be aborted or timed out) unless they are matched by a Restriction rule, in which case they will not be touched by Crash Protection under any circumstances as defined on the CP Restrictions page.

What is a Survival Strategy and Self-healing?

The survival strategy is how FusionReactor will attempt to self-heal the situation if crash protection is triggered. It can be configured to kill problem requests and queue new requests until the situation is stabilized. It will also alert you so that you have time to resolve the problem if you need to take action.

Survival Strategy:

Configuration	Description
Abort Strategy	When a request is aborted due to Crash Protection there are two possible behaviours: • **Display Message** - Enter your message in the text box provided and that message will be presented to the user who made the aborted request. • **Redirect to URL**- Enter a URL in the text box provided and the user who made the aborted request will instead be redirected to this URL. **NOTE** Redirects may not work when a page has already sent data to the browser.
Add Parameters	After an abort, if you "Redirect to URL" then this option defines whether FusionReactor will add extra parameters to your URL so you can see exactly what happened. If you enable parameters the URL will include the following: • **detection_method** - This is the type of Crash Protection which fired. It will be one of: "freemem", "requests", or "timeout" • **threshold_value** - This is the threshold value associated with the type of Crash Protection, as defined earlier on this page. • **actual_value** - This is the actual value which caused the Crash Protection to fire. • **source_url** - Contains the URL sent on the original request e.g. http://127.0.0.1/test.cfm (since 3.0.1) • **source_url_params** - Contains the URL parameters sent on the original request (since 3.0.1) • **source_method** - This is the HTTP method used on the original request e.g. GET, POST (since 3.0.1) • **start_time** - This is the start time of the request from Epoch e.g. 1209376581900 (since 3.0.1)
Queue Timeout / (Request Abort) (seconds)	If a request is queued due to Crash Protection, this value represents the longest amount of time that the request will be held before FusionReactor abandons it.The default value is 60 seconds.

Crash Protection Logging:

Configuration	Description

CP Logging	If Crash Protection Logging is enabled then Crash Protection information will be written to the Crash Protection Log. If Disabled is chosen then Crash Protection information will not be written to the Crash Protection Log.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the Crash Protection Log from inside FusionReactor.

How do I set up Notifications?

From the FusionReactor Administrator, click "Settings" from the table of contents (FusionReactor section) and you will see the FusionReactor Settings screen. The Email section of this screen deals with the notification.

Why would I want to enable Crash Protection?

There are many cases in which crash protection can be an invaluable tool. These include:

Easing Support

Crash Protection automates the manual process of checking each of your servers and verifying that they are operating within satisfactory parameters. Instead of having to periodically go through each server you can have warnings sent directly to your inbox.

Request Bursts

If you have a server which receives bursts of requests all at once then FusionReactor can using queuing spread that load out over time and so protect your server from becoming overwhelmed.

Preventing Crashes

If you periodically have very resource hungry jobs running then FusionReactor can stop new requests being run which could potentially cause the server to become unstable.

During Development

If you need to write a page which cannot use more that a fixed amount of the system resources then crash protection is a good way of tracking where and when you have problems. As well as emails, crash triggers can also be viewed in a log file.

Crash Protection Log (crashprotection.log)

The Crash Protection Logs are a set of rotating files which you can configure from the Crash Protection Settings page.

The log file is formatted as follows:

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time at which this log entry was written.

Date/Time (ms)	3 [C]	The number of milliseconds since the epoch.
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
CP Type	6 [F]	The type of Crash Protection which caused this log entry to be written.
Action	7 [G]	The action which was carried out on this request by Crash Protection.
Memory Percentage	8 [H]	The percentage of free memory at the time of this log entry.
Memory Threshold	9 [I]	The threshold percentage of free memory set up in Crash Protection settings.
Request Count	10[J]	The number of currently running requests at the time of this log entry.
Request Threshold	11 [K]	The threshold number of currently running requests set up in Crash Protection settings.
Execution Time	12 [L]	The current execution time of this request at the time of this log entry.
Execution Threshold	13 [M]	The threshold execution time set up in Crash Protection settings.
Server ID	14 [N]	The ID of the server for Server Up / Down notifications
Script	15 [O]	The script that you set up in the Crash Protection settings (including passed parameters if available).
Date (Formatted)	16 [P]	The date on which this request was started or finished (depending upon the Request Status.)

Time (Formatted)	17 [Q]	The time at which this request was started or finished (depending upon the Request Status.)
Date/Time (ms)	18 [R]	This is the millisecond time at which this request was started or finished (depending upon the Request Status.)
Version	19 [S]	Version number of this log entry.
Server Startup Time (ms)	20 [T]	Startup time of the application server in milliseconds (ms) since epoch.
Request ID	21 [U]	The FusionReactor Request ID.
Request Status	22 [V]	The current state of the request. This will tell you if a request has started, finished, was killed, queued, etc...
CP Reason	23 [W]	The reason given by Crash Protection for the current action.
Thread ID	24 [X]	The name of the thread responsible for responding to this request.
Client IP Address	25 [Y]	The IP address of the machine making the request.
Request Method	26 [Z]	This will usually be "GET" or "POST"
Request URL	27 [AA]	This is the requested URL. The format of this column will change depending on URL Format key which can be found on the Request Settings page.
Execution Time (ms)	28 [AB]	The amount of milliseconds it took to complete the request. (For incomplete requests, this column will be 0.)
Used Memory (percentage)	29 [AC]	The amount of memory (expressed as a percentage) which was used when this request started or finished.
Max Memory (KB)	30 [AD]	The total amount of physical memory available to this instance.

Used Memory (KB)	31 [AE]	The amount of memory which was used when this request started or finished.
Total Memory (KB)	32 [AF]	The amount of memory which was allocated by the instance when this request started or finished.
Free Memory (KB)	33 [AG]	The amount of free memory (within the allocated block) when this request started or finished.
Query String	34 [AH]	If the URL has any parameters then they will appear here.
Return Status Code	35 [AI]	This is a HTTP return code such as 200 (OK,) 404 (Not found,) or 500 (Internal Server Error.) For "Started:" rows this column will be 200.
CPU Time (ms)	36 [AJ]	This is the amount of actual CPU time which this request required. (For incomplete requests, this column will be 0.)
AMF Request	37 [AK]	If AMF decoding is enabled (Request Settings page) and there is AMF to decode, then the method names will appear here.
JSESSIONID	38 [AL]	The J2EE Session Id for this request.
CFID	39 [AM]	The CF Id for this request.
CFTOKEN	40 [AN]	The CFTOKEN for this request.
JDBC Query Count	41 [AO]	The number of JDBC queries run by this request at the time of this log.
JDBC Total Time (ms)	42 [AP]	The total amount of time spent running JDBC queries at the time of this log.
JDBC Total Execution Time (ms)	43 [AQ]	The total amount of time spent by the database running JDBC queries at the time of this log.
JDBC Total Rows	44 [AR]	The total number of rows returned by JDBC queries at the time of this log.

Bytes Sent	45 [AS]	The amount of data which was sent back to the client.
Time to First Byte (ms)	46 [AT]	The number of milliseconds it took to deliver the first bit of data.
Time to Last Byte (ms)	47 [AU]	The number of milliseconds it took to deliver the complete content.
Time to Stream Open (ms)	48 [AV]	The number of milliseconds before the data stream was opened.
Time to Stream Close (ms)	49 [AW]	The number of milliseconds until the completed data stream was closed.

Deleting the Crash Protection Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the Crash Protection Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See

Log Settings
Importing and Graphing Data in Excel

JDBC

JDBC Settings

JDBC Stack Trace Filter

JDBC Log (jdbc.log)

JDBC Settings

The JDBC Settings page lets you modify various elements relating to JDBC. The page has the following items:

SQL Statement Recording (Request Detail Page)

Configuration	Description
Record SQL	Specifies how SQL is stored along with each request. You can either store queries in slowest first order (Slowest:default) or in the order that they are run (In Order). Alternatively you can just get a summary of run queries (Summary Only).

Record up to (queries)	This limits the number of queries stored for each request. If you are storing them in order then you will see the first queries which are run. The number of SQL statements recorded on the request affects the amount of memory that FusionReactor could use. It is recommended that you increase this value in stages and check the memory usage after each increase once the server has run at least the number of request that are kept in the Request History. Each query recorded only uses a very small amount of memory but this may become significant if you store hundreds or thousands of queries per request when multiplied by the number of requests maintained in your request history setting. When queries are recorded in Slowest Order and the number of queries to record has been reached, a slower query will be correctly inserted into the list of recorded queries according to the time it has taken to run, and the last entry in the list will removed and added to the Summary Information.
Only queries slower than (ms)	This threshold value lets you record only SQL statements that run slower than the amount of time you specify here. If the SQL statement runs faster than the time you specify it's data will still be added to the Summary information.

Additional Information

Configuration	Description
Query Location	When this feature is enabled and a JDBC request is executed, FusionReactor will store a Java Stack Trace which can be used to find out the location of each query that ran. You change the amount of information which gets stored in this stack trace by going to the JDBC Stack Trace Filter page. The resulting stack traces are visible from the Request Detail Page page and within the JDBC Log Files.

JDBC Logging (Log Files)

Configuration	Description
JDBC Logging	If JDBC Logging is enabled then JDBC information will be written to the JDBC log.
Only queries slower than: (ms)	Only log queries that run longer than this time will be written to the log files. (Optional)

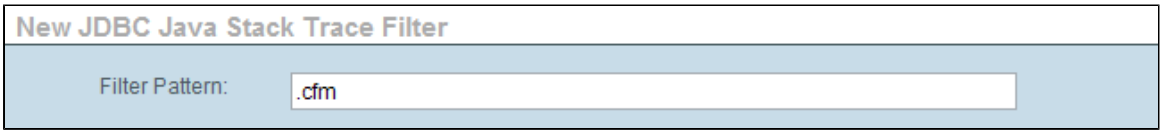
View Size (KB)	This value indicates how much of the log file should be displayed when you view the JDBC log from inside FusionReactor.

> **NOTE**
>
> These settings pertain only to DSNs wrapped with the FusionReactor JDBC Driver Wrapper. For more information, please see the JDBC Driver Wrapper User Guide.

JDBC Stack Trace Filter

The JDBC Stack Trace Filter page lets you set up various filters which are used to make the Java Stack Trace more manageable. The complete Java Stack Trace contains far more information than you are likely to need, so this page has been added to let you tailor it to your own preferences. By default, FusionReactor will show you any Stack Trace rows which contain references to cfm, cfc or jsp pages.

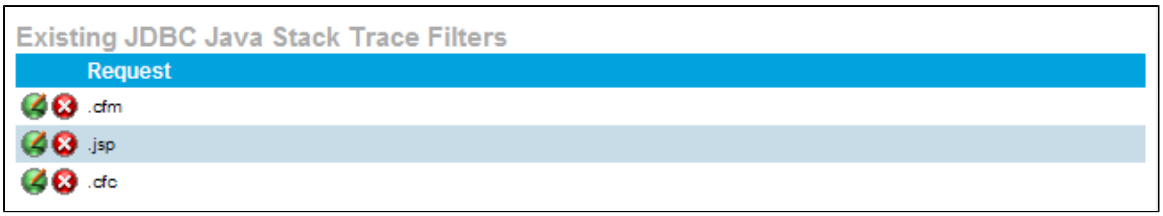
To add a new filter, just type a string into the field provided and click the "Add Filter Pattern" button. Any lines in the stack trace which contain that string will be stored along with the JDBC request.



Managing Existing Filters

Below the Add Filter Pattern form you will see a list of all existing filter patterns. Each pattern has two icons next to it:

Button	Name	Description
	Modify Pattern	Clicking on this icon will take you to the Modify JDBC Java Stack Trace Filter page. From this page you can change the filter pattern (Described in the paragraph at the top of the page).
	Delete Pattern	Clicking on this icon will prompt you if you are sure you want to delete this filter pattern. If you OK the dialog then the filter pattern will be removed from the Existing JDBC Java Stack Trace Filters list and the page will be refreshed.



JDBC Log (jdbc.log)

The JDBC Logs (which are a set of rotating files which you can configure from the JDBC Settings page) are simple space separated files, so they can be easily imported into database or spreadsheet applications. The log is called **jdbc.log**, the bottom part of the page tells you where you can locate this file on your hard drive.

The log file is formatted as follows:

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date/Time (Milliseconds)	3 [C]	This is the millisecond time at which this request was started or finished (depending upon the Request Status.)
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
Request ID	6 [F]	The FusionReactor Request ID.
Thread ID	7 [G]	The name of the thread responsible for responding to this request.
Client IP Address	8 [H]	The IP address of the person making the request.
Request Method	9 [I]	Specifies the HTTP method used during the request. This will usually be "GET" or "POST"
Request URL	10 [J]	This is the requested URL. The format of this column will change depending on URL Format setting which can be found on the Request Settings page.

Log Type	11 [K]	One of: "METRIC", "NOTIFICATION" or "REMINDER"
DB Start Time	12 [L]	The time (in milliseconds) at which the JDBC query began.
DB End Time	13 [M]	The time (in milliseconds) at which the DB finished with the query.
Total End Time	14 [N]	The time (in milliseconds) at which the query was completely finished (after results were transferred back etc.).
DB Time	15 [O]	The amount of time (in milliseconds) this JDBC query spent in the DB.
Total Time	16 [P]	The total amount of time (in milliseconds) it took to complete this query.
Row Count	17 [Q]	The number of rows returned by this query.
Prepared Statement	18 [R]	This is a simple TRUE/FALSE value to tell you if this statement had already been prepared prior to its execution.
Row Limited	19 [S]	This is a simple TRUE/FALSE value to tell you if this query was row limited.
DataSource	20 [T]	If you have assigned a name to your data source wrapper then it will appear here. To find out how to name your data sources, please refer to page 14 of the JDBC Driver Wrapper User Guide.
Statement	21 [U]	The actual statement which was run.
Stack Trace	22 [V]	This is the stack trace which was stored when the query was run. You can turn on or off this feature from the JDBC Settings page and alter the stack trace information from the JDBC Stack Trace Filter page.

Query String	23 [W]	This is the query string from the request running the JDBC statement.
Message	23 [X]	For Notifications and Reminders, this will contain the notification or reminder message.

> **NOTE**
>
> Depending on the type of log message, not all columns will have values.

Deleting the JDBC Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the JDBC Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See

Log Settings
Importing and Graphing Data in Excel

Compression

Compression Settings

MIME Type Restrictions

Exclude URLs

Compression Settings

The Compression Settings page lets you turn on or off the compression features as well as letting you define how the MIME Type Restrictions behave.

Compression Support:

Configuration	Description
GZIP Requests	When this option is enabled, FusionReactor will attempt to decompress HTTP requests which are sent to the server in a compressed format.

GZIP Responses	When this option is enabled, FusionReactor will attempt to compress outgoing HTTP responses unless otherwise instructed by the MIME Type Restrictions or the URL Exclusions. **NOTE** FusionReactor will only compress responses when the client request indicates that the client can support compression by sending a_ Accept-Encoding: gzip *header*

MIME Type Restrictions:

Configuration	Description
MIME Type Restrictions	This value defines how MIME Type Restrictions behave. • Ignore requests that match the rules - By default, all requests will be compressed when compression is enabled unless the page's document mime type matches one of the rules defined on the MIME Type Restrictions page (then it will **NOT** be compressed). • Compress requests that match the rules - no requests will be compressed unless the page's document mime type matches one of the rules defined on the MIME Type Restrictions page (then it will be compressed).

MIME Type Restrictions

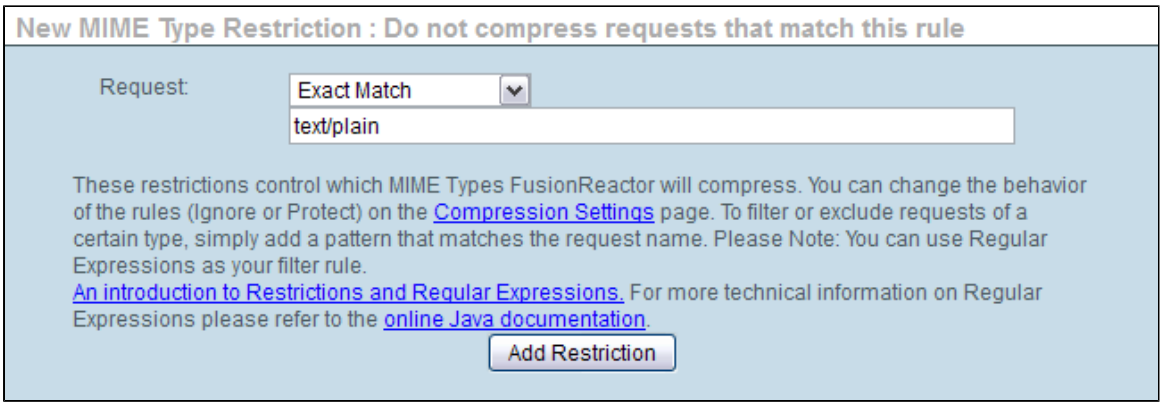
MIME Type Restrictions are designed to let you specify which types of file you want to be compressed, or alternatively, which types you do NOT wish to be compressed (leaving all other file types to be compressed). To select which mode the Restrictions page works in, change the Restrictions option on the Compression Settings p age. The MIME Type Restrictions page has two main areas: At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for this feature.

New MIME Type Restriction

To enter a restriction you will need to fill the following fields:

Configuration	Description

Request	This is the document type that you want to restrict. Firstly select if you are creating an "Exact Match" or "Regular Expression". Next enter your MIME Type. For examples of Regular Expressions, please see Regular Expressions and Content Filters.
Text or Regular Expression	The text input here depends on the option you choose for Request. The options are "Exact Match" or "Regular Expression".



Once you submit the form, the page will be refreshed and your new restriction will be listed at the bottom of the page.

Managing Existing Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

Button	Name	Description
	Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above (New MIME Type Restriction).
	Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.



See Also:

Compression Settings

Exclude URLs

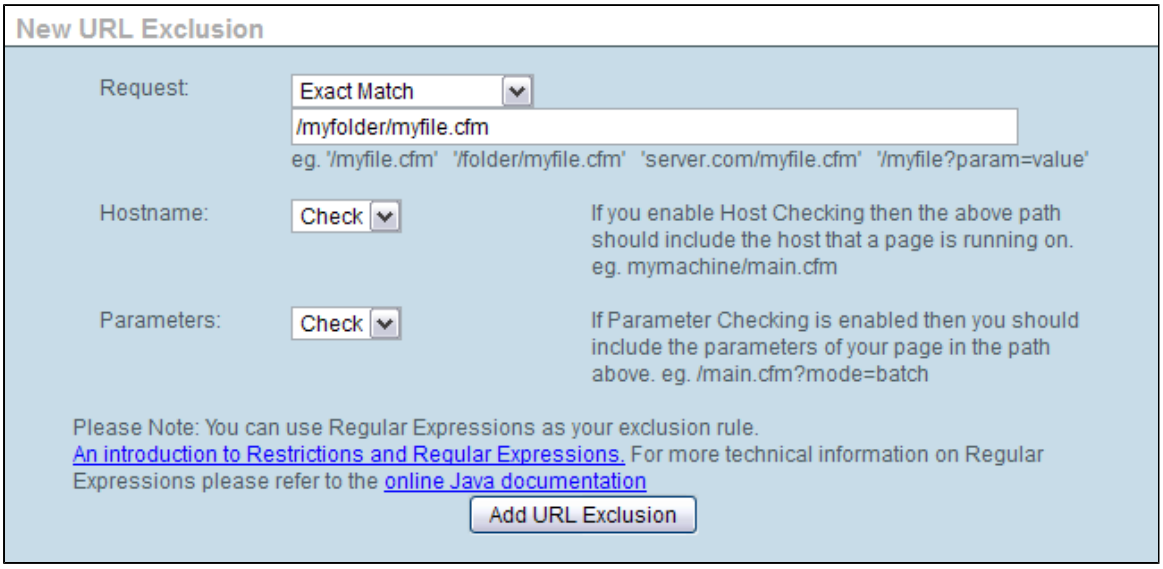
The Exclude URLs list lets you define a set of pages which will never be compressed. The Exclude URLs page has two main areas: At the top of the page there is a form for you to enter new exclusions and below this is a list of all existing exclusions for this feature which you can edit or delete.

New URL Exclusion

To enter an exclusions you will need to fill the following fields:

Configuration	Description
Request	This is the page you want to exclude. Firstly select if you are creating an "Exact Match" or "Regular Expression". **Exact Match** Next enter the location of your page. By default the request will be an exact match and should look something like "/batchStats.cfm" or "/batch/dailyStats.cfm" (beginning with a '/' and without parameters.) . The text you input here will depend on the options you choose for the Hostname and Parameters fields. **Regular Expressions** A good use of Regular Expressions here would be to restrict everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Restrictions. The text you input here will depend on the options you choose for the Hostname and Parameters fields.
Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you enable Hostname checking then the request field should also include a hostname and look something like "productionMachine/batchStats.cfm". You should NOT include HTTP or HTTPS at the start of this string.

Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example "/batchStats.cfm?period=YEAR". Once again, Regular Expressions and Restrictions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.



Once you submit the form the page will be refreshed and your new exclusion will be listed at the bottom of the page.

Managing Existing Exclusions

Below the Add Exclusion form you will see a list of all existing exclusions. Each exclusion has two icons next to it:

Button	Name	Description
	Modify Exclusion	Clicking on this icon will take you to the Modify Exclusion page. From this page you can change any of the information which is listed above (New URL Exclusion).
	Delete Exclusion	Clicking on this icon will prompt you if you are sure you want to delete this exclusion . If you OK the dialog then the exclusion will be removed from the Existing Exclusion list and the page will be refreshed.



Content Filters

Search and Replace

Content Filter Restrictions

Filter Settings

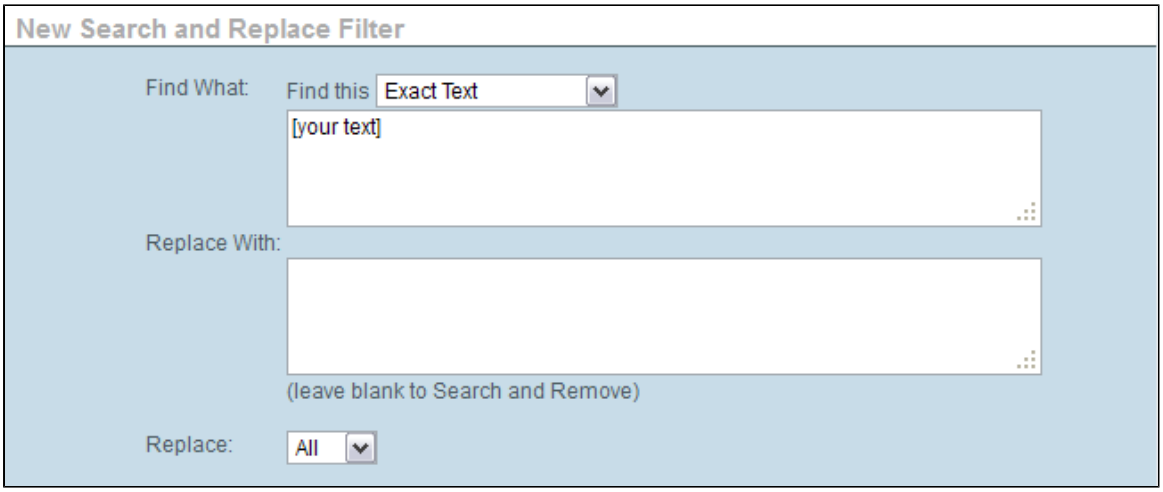
Search and Replace

The Search and Replace feature is arguably one of the most powerful features in FusionReactor. It allows you to make temporary changes to your pages without having to modify source code. Fix broken links, spelling mistakes or html problems. Search and Replace Filters enable you to change the content sent to the web browser by replacing or removing text on the fly. The Search and Replace page has two main areas: At the top of the page there is a form for you to enter new filters and below this is a list of all existing filters for this feature.

New Search and Replace Filter

To enter a Search and Replace filter you will need to fill the following fields:

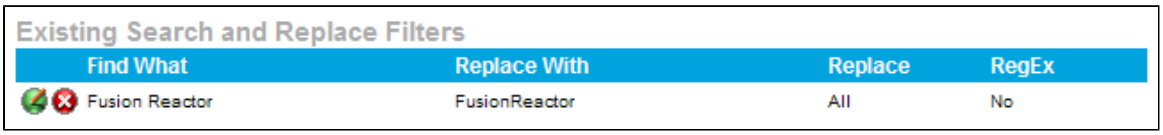
Configuration	Description
Find What	Firstly you should select if you want to search for an Exact String, or if you want to use Regular Expressions to search for a pattern. For examples of Regular expressions, see Regular Expressions and Content Filters. Once you have selected the search type, you should enter the your search string in the text field provided.
Replace With	This field contains the text you want to replace your search string with. If you just want to search for a string and remove it from your page then simply leave this field blank.
Replace	This option lets you decide if you want to replace all instances of your search string (All), or just the first one found on the page (First).



Managing Existing Search and Replace Filters

Below the Add Group form you will see a list of all existing filters. Each group has two icons next to it:

Button	Name	Description
	Modify Group	Click this icon to change a group name or to change the servers which are currently a part of this group (New Search and Replace Filter).
	Delete Group	Clicking on this icon will prompt you if you are sure. OK the dialog box to delete the group.



Content Filter Restrictions

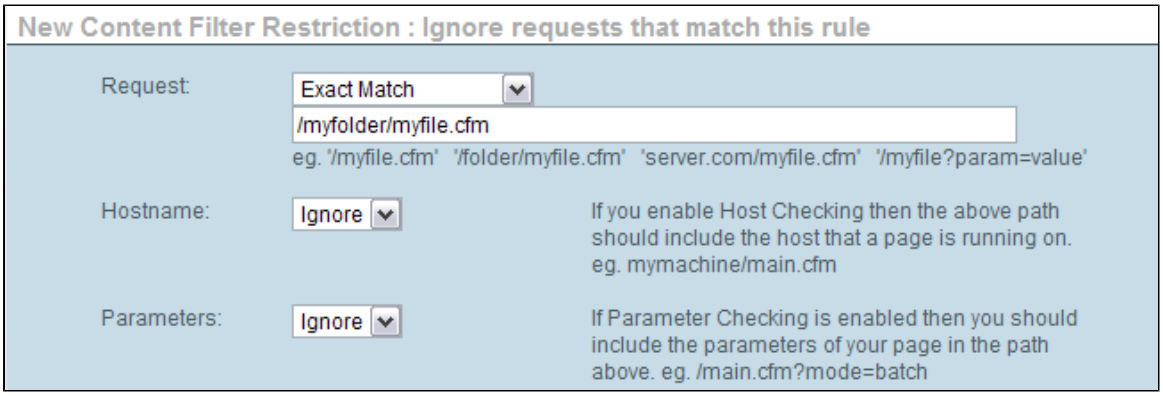
Content Filter restrictions are designed to let you specify which pages you want run through the content filter system, or alternatively, which pages you do NOT wish to be filtered (leaving all remaining pages to run through content filtering.) To select which mode the Restrictions page works in, change the Restrictions option on the Content Filter Settings page. The Content Filter Restrictions page has two main areas: At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for this feature.

New Content Filter Restriction

To enter a restriction you will need to fill the following fields:

Configuration	Description
Request	This is the page you want to exclude. Firstly select if you are creating an "Exact Match" or "Regular Expression". **Exact Match** Next enter the location of your page. By default the request will be an exact match and should look something like "/batchStats.cfm" or "/batch/dailyStats.cfm" (beginning with a '/' and without parameters.) . The text you input here will depend on the options you choose for the Hostname and Parameters fields. **Regular Expressions** A good use of Regular Expressions here would be to restrict everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Content Filters. The text you input here will depend on the options you choose for the Hostname and Parameters fields.

Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you enable Hostname checking then the request field should also include a hostname and look something like "productionMachine/batchStats.cfm". You should NOT include HTTP or HTTPS at the start of this string.
Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the Hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example "/batchStats.cfm?period=YEAR". Once again, Regular Expressions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.



Once you submit the form the page will be refreshed and your new restriction will be listed at the bottom of the page.

Managing Content Filter Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

Button	Name	Description
	Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above (New Content Filter Restrictions).

	Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.



Filter Settings

The Content Filter Settings page lets you turn on or off the Search and Replace feature as well as letting you define how the Filter Restrictions behave.

Content Filters:

Configuration	Description
Search and Replace	When this option is Enabled, FusionReactor will run responses through the Search and Replace feature unless otherwise instructed by the Filter Restrictions.

Content Filter Restrictions:

Configuration	Description
Filter Restrictions	This value defines how Filter Restrictions behave. • **Filter requests that match the rules** - By default, no requests will be run through the Search and Replace filters. If a request matches one of the rules defined on the Filter Restrictions page then it will be filtered. • **Ignore requests that match the rules** - By default, all requests will be run through the Search and Replace filters. If a request matches one of the rules defined on the Filter Restrictions page then it will not be filtered.

Plugins

Active Plugins

Open Services Gateway initiative (OSGi)

Reactor Log (reactor.log)

Active Plugins

The table lists all plugins that are active and currently running on FusionReactor.

Name	Description
Plugin Information	This contains a brief description of the plugin as well as the name of the jar file that is the plugin.
Version	The version of the plugin as well as the name of the plugin.
Type	Lists the bundle framework (OSGi) and the type of plugin.
Actions	**Logfile:** Contains information about when the plugin was started up and shutdown. **Configuration:** Clicking this allows you to configure the plugin. Refer to the individual plugin definitions for details regarding Plugin configuration

Also See

Overview of FusionReactor Logs

FusionReactor Plugins

Open Services Gateway initiative (OSGi)

Open Services Gateway initiative (OSGi)

FusionReactor 4.0 introduces a new Open Services Gateway initiative (OSGi) based plugin architecture to enable new features to be seamlessly deployed. This architecture allows us to deploy plugins (coming in the form of bundles for deployment), which can be remotely installed, started, stopped, updated and uninstalled, without the necessity for a reboot.

OSGi Bundle Inventory

Name	Description
Bundle Information	Contains a link to a website which contains further information about the bundle, a brief description of the plugin and the name of the jar file that is the bundle.
Version	Contains the version of the bundle as well as the name of the bundle.
ID	Contains the ID assigned to each bundle, the number can vary from execution to execution.
Type	Lists if the software is a bundle or a bundle and a plugin.

Start Level	Each bundle is executed starting from the lowest value to the highest.
Status	If the status reads ACTIVE then the plugin is running, if it reads RESOLVED it isn't running however it has the necessary files to be able to run.
Actions	**Stop:** Stops execution of the plugin. **Start:** Starts execution of the plugin. **Logfile:** Contains information about when the plugin was started up and shutdown. **Configuration:** Clicking this allows you to configure the plugin, see the FusionReactor Plugins page for a guide on different plugin configurations.

OSGi Environment

Name	Description
Framework	The framework that the OSGi uses.
Started	The date when the OSGi environment began execution.
Current Run Level	The current run level.
Properties	Various properties about the OSGi environment, this is mostly used for debugging. **NOTE** When technical support is required this information may be requested from you.

Status Log - osgi.log

As OSGi is responsible for the deployment of plugins this log contains information relating to that as well as the state of OSGi itself. Below is a table describing what all the attributes mean in the log file (**osgi.log**).

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.

Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.

Also See

Log Settings
FusionReactor Plugins

Reactor Log (reactor.log)

A log will be generated (reactor.log). The log contains information about FusionReactor as whole, typical it contains information on settings that have been changed as well as information critical to it. Below is a table explaining what all the attributes in the log file mean.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.

Also See

Log Settings
FusionReactor Plugins

Logging

This section contains information on logs archives and log settings. Both are pages you can visit in

FusionReactor by clicking Logging. Both pages talk about the settings you can configure in those pages within FusionReactor.

Log Archive

Log Settings

Overview of FusionReactor Logs

Log Archive

This page is used to view information on the log archives that are generated by FusionReactor as well as perform some functionality on the logs.

Table Column Headings

Name	Description
Archive File	The name of the archive file stored on disk.
Size	The size of the archive file (measured in megabytes)
Created	The date when the archive file was created.
Source	The name of the server that the logs are based on.
Archive Comment	If the text reads Don't Keep then that archive will be removed after a purge has taken place. If it reads Keep then the archive will remain on the drive after a purge has taken place. Simply click the text to change between the two.
Keep	If the text reads Don't Keep the archived will be removed from disk after the purge, if it reads Keep it will remain on disk after the purge has taken place. To change it click on the text and will will change to the other option.

Table Buttons

Each row contains the following two buttons on the left hand side, the effects of clicking them is described below.

Button	Name	Description
	Delete Archive	Upon clicking this you will be asked if you remove the archive, choosing yes will remove it from FusionReactor and the machine that it was stored on.

	Show History	Clicking this will display additional information about the archive. The page is split into two different sections, one contains information about the archive file and the other contains information about how the file was created. Below is two lists summarizing the key points from both sections of the page. **Top section:** • **Archive Filename :** The absolute file path of the archive. • **Original File:** The name of the archive file. • **Timezone:** The timezone used by FusionReactor. • **FusionAnalytics Target:** Specif ies if the archive is to be used by FusionAnalytics if so the FusionAnalytics target is listed. **Bottom Section:** • **Date :** Date the archive was created. • **Source System :** The software used to capture the files for the archive and to create it. • **Subsystem :** The software component within the Source System (FusionReactor) used to perform an action. • **Message :** This is a brief description of the action that the Subsystem performed.

Rotate Logs

To Rotate the logs click on the button that says Rotate Now. This will generate a new archive based on the log information gathered since the previous archive was generated. The archive will then be written to disk.

Purge Now

To purge the logs click the button that says Purge Now. Selecting this will purge (remove) archives based on rules that have be created by the user. Information about purging can be found in the Log Settings section.

Clear All Archives

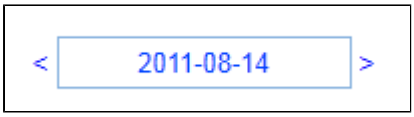
To clear all the archives click on the button that says Clear All Archives. Clicking this will ask you if you want to remove all archives, clicking yes will remove them from the system.

WARNING

Make sure you are 100% sure before removing all archives as doing so may compromise your system.

View Logs From a Different Date

Clicking the "<" symbol allows you to revert back to a previous date allowing you to view the details of older logs, clicking the ">" symbol advances the date allowing you to move through time and view more recent logs.



Log Settings

This page is used to alter log engine settings, log rotation settings, lot archive settings and log archive monitor settings. The page is broken up into sections, each containing a table explaining the options that can be chosen in drop down boxes or input into text fields.

Locating Log Files

The location of the FusionReactor logs will depend on the mode of logging you are using and how you have configured your Log Directory. Below is a table covering the two types of log rotations that exist in FusionReactor.

Rotation Mode	Description
Centralized Archive and Rotation mode	The current log directory will be in a time-stamped folder beneath the path specified in Log Directory. `FusionReactor/Instance/your_instance_name/log/time-stamped` folder. In order to have logs zipped and archived automatically, and to ship them using the **FusionAnalytics Connector** to a FusionAnalytics instance, you **must** choose this option.
Traditional Per-Log Rotation mode	Logs are written directly to the path specified in Log Directory. Changing this setting will cause new log entries to be written to the new location. You can find FusionReactor's logs in the `/FusionReactor/Instance/your_instance_name/log` directory. These are rotating files, the newest having the suffix -0.log, the next oldest -1.log etc.

Log Rotation Settings

Configuration	Description

Rotation Mode	**Periodic** This will update the Log Archive every *n* minutes, where *n* can be specified in the text field below. **Fixed** This will update the Log Archive daily using a time that can be specified in the text field below (HH:MM in the 24 hour system).
Rotate Every/At	If you choose **Periodic** as the rotation mode, then you input here how often you want to rotate the log (in minutes). If you choose **Fixed** then enter here at what time of the day you want the logs to be rotated (HH:MM in the 24 hour system).
Rotate After Crash	Choosing Enabled will rotate the logs after a crash has occurred, also if FusionAnalytics is available it will queue them for transmission. Choosing Disabled will mean that the logs aren't rotated after a crash.

Log Archive Settings

Configuration	Description
Retention Strategy	You can choose to keep archives for a certain number of days, or until they reach a certain size (MB).
Retention Value	Input a value based on the Retention Strategy you choose. If you choose to keep them for a certain number of days, this value is that number. If you wish to retain only a certain total size of archive, enter the size (in MB) here.

Automatic Purge	If you set this option to Enabled then archives will be purged based on the Retention Strategy you choose, if you choose Disabled then archives will never be purged. **WARNING** Disabling this option will allow archives to grow without bounds however enabling it will cause unsecured (not backed up) archives to be removed. Carefully consider what option you select here. The archive is purged when you select the **Purge Now** button in the Connector Status or Log Status page, and after a new archive zip is added to the archive. The purge mechanism always considers oldest archives first, deleting until the **retention strategy** sel ect above is satisfied. Archives marked "**Keep**" (in Connector Status or Log Status) are never considered for purge.
Set "Don't Keep" by default	Sets the **Keep/Don't Keep** option on new archives. When enabled, archives are immediately eligible for purge, regardless of whether they have been shipped automatically to FusionAnalytics. When disabled, archives are not eligible for purge until a successful transfer to FusionAnalytics has occurred.
Archive Location	The location of where the archives reside on disk is input here. A default location is provided.

Log Archive Monitor Settings

Configuration	Description
Log Archive Mail Job	If set to enabled, the user will receive an email daily about the disk usage of the log and archive system. See Settings for help configuring the email server.
Scheduled Time	If the above option is enabled, input the time (HH:MM in the 24-hour system) the monitor job should run.
Warning Threshold	If the total archive size reaches the input number here (megabytes) then an email will be sent warning the user of the excessive amount of space being used. The title of the email will read *[WARNING!]*

Overview of FusionReactor Logs

FusionReactor provides many helpful graphs and metrics to help you see how your system is performing. These graphs are available in real-time and will also show you a short history of what happened (depending on the amount of data which has been generated), but what happens when your server begins to crash? Lots of this data is stored in-memory so it can be used to generate reports and lists for the FusionReactor Administrator. If you are having problems which cause (or necessitate) restarts of the server process, this in-memory data is lost - however all of this data will have been captured to one of FusionReactor's extensive log files.

This data can in almost all cases be used in a post-mortem session to pinpoint the location of problems within a system. In addition to immediate post-mortem work, the logs can also be used to analyze usage patterns, trends and to derive longer-term aggregate data.

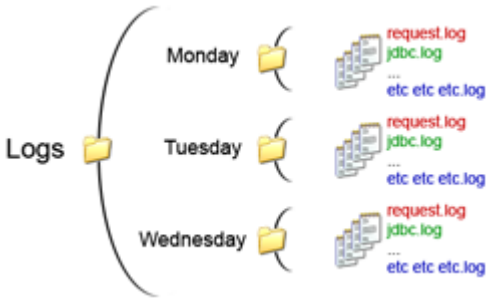
As of FusionReactor 4.0, there are two styles of logging which is supported by FusionReactor. You can change current logging method on the Log Settings page.

> ⚠ If you wish to use the FusionAnalytics Connector then you must be using the Centralized Archive and Rotation logging method. If you try to enable the FusionAnalytics Connector whilst in Traditional mode then you will be be notified that the logging method will need to change and be given the option of having FusionReactor automatically make that change for you.

Centralized Archive and Rotation

If you install FusionReactor from scratch then this is the logging mechanism you will be using by default. The idea behind Centralized Archive and Rotation is that you end up with a set of log folders named with a timestamp (within the existing Log Directory specified on the Log Settings page) that each contain a complete set of log files for a specified period of time. You can choose whether a new folder be created every 'n' minutes, or you can specify a time at which the logs are daily rotated. Additionally you can specify how much log data you want to save, either by amount of disk space used, or period of time covered (or you can save all the logs if you have a big enough hard drive). The main advantage to this logging method over the older Traditional method is that is is a lot easier to batch up a complete set of log files covering a distinct period of time which is very useful when trying to compare values between log files or when importing them into something like FusionAnalytics.

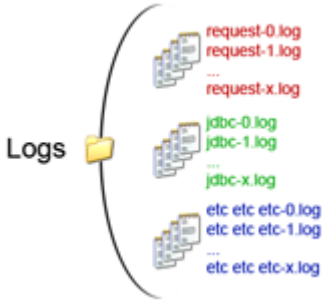
All of these log settings are available on the Log Settings page.



Traditional Per-Log Rotation

If you upgrade an older version of FusionReactor to 4.0 then this will be the default logging mechanism you will be using. It is the logging mechanism used by all versions of FusionReactor prior to 4.0. The idea behind this logging mechanism is that, for each log file in the system, you can specify how big that file is allowed to become and how many log files you want to keep. If you specify a log size of 10MB and say that you want a history of 5 log files then you will have a maximum of 50MB of logs available to you at any time. With Traditional Per-Log Rotation, all the log files are stored in the Log Directory specified on the Log Settings page. For each type of log (eg request.log) you will end up with a set of numbered log files (eg. request-0.log, request-1.log ... request-4.log). Once request-0.log is full, all the log files are renamed and a new request-0.log is started (so request-0.log would always contain the latest log messages). The oldest log file (in this example, request-4.log) is deleted and immediately replaced with the file previously named request-3.log.

The settings for the File Size and File Count of your log files can be found in the associated settings page (eg. Request Settings for the request.log). We recommend however that you switch to the new Centralized Archive and Rotation logging method as it will give you more control over your logs.



Log files available in FusionReactor

Below you will find an overview of all of the log files which are available in FusionReactor, with a short description of each file. Click on the log file name in the table to see a detailed description of all the log columns. You can also find a **[*TYPE*]-headers.txt** file in the log folder, for the main log files. This file contains the symbols used in the logs.

Data Logs

These are logs that relate to captured server data, for example SQL results, memory information etc.Clicking on the log names will take you to a page dedicated to that log, it will explain each attribute in the log as well as provide you with useful additional information.

Log Name	Log Type	Created By
request.log	DATA	SYSTEM
resource.log	DATA	SYSTEM
jdbc.log	DATA	SYSTEM
crashprotection.log	DATA	SYSTEM
stacktrace.log	DATA	SYSTEM
heartbeat.log	DATA	PLUGIN
deadlock.log	DATA	PLUGIN
dbpoolstats.log	DATA	PLUGIN
hitcountstats.log	DATA	PLUGIN
realtimestats.log	DATA	PLUGIN
memorysummary.log	DATA	SYSTEM
[memory_name].log	DATA	SYSTEM
cputhread.log	DATA	SYSTEM

fr-jrun-event.log	DATA	PLUGIN
reactor.log	STATUS	SYSTEM
notification.log	STATUS	SYSTEM
osgi.log	STATUS	SYSTEM
plugin-fr-amf-decorator-plugin.log	STATUS	PLUGIN
plugin-fr-fa-connector-plugin.log	STATUS	PLUGIN
plugin-fr-heartbeat-plugin.log	STATUS	PLUGIN
plugin-fr-logrotator-plugin.log	STATUS	PLUGIN
plugin-fr-jrun-plugin.log	STATUS	PLUGIN
plugin-fr-deadlock-plugin.log	STATUS	PLUGIN

Key

Word	Definition
DATA	Data related to the server memory, CPU, response time etc.
STATUS	Data related to a PLUGIN or the SYSTEM. The information is brief and only covers the state of the PLUGIN or SYSTEM.
SYSTEM	This is FusionReactor.
PLUGIN	A plugin that is connected to FusionReactor.

Log File Format

All log spaces are space deliminated. This allows you to easily import the files into spread sheets such as Microsoft excel and openoffice as well as databases, this allows you to view the logs how you want. Below is an example of a log file as well as the table that would accompany it.

Example log file and table

Below is an example log file

SERVER1 70 5 21

This log file contains the attributes CPU, 100, 5 and 21. Below is the table that would explain the log file.

Field Name	Number	Description

Server Name	1 [A]	The servers name.
Efficiency(percentage)	2 [B]	How efficient the server is.
Downtime (hours)	3 [C]	The downtime of the server.
Uptime (hours)	4 [D]	The up time of the server.

Exceptions to the rule

There are several log files which also contain a plain text message. (eg reactor.log) In this case, the normal rule applies for all the values up to the last 'message' value, but after that, everything up to the end of the line is considered to be the value of the last field.

Also See

Analytics

Importing and Graphing Data in Excel

Methodology in OpenOffice, where we used the **request log** with standard deviation to derive a nominal page run time value.

Analytics

The Analytics menu provides access to the configuration of the FusionAnalytics Connector, a system which provides automatic log shipping between FusionReactor instances and FusionAnalytics.

FusionReactor logs very detailed metrics on various aspects of a J2EE container: request runtimes, memory loads, CPU occupancy, JDBC activity and CPU-Thread usage, to name but a few. FusionReactor derives immediate value from this data by way of sophisticated charting and tables. Getting a longer perspective - to allow resource planning, trend discovery etc. - is outside the scope of FusionReactor, but is well suited to the **FusionAnalytics** product.

FusionReactor, using the FusionAnalytics Connector, is able to ship logs automatically to a FusionAnalytics DataCollector, where it can be easily stored in a database and made available for a comprehensive, extensive set of analytics operations.

More information on FusionAnalytics is available here.

Connector Status

FusionAnalytics Settings

FusionAnalytics DataCollector Targets

Locked

Naming Applications and DataCollectors

FRAM Target

Connector Status

This page contains a table listing all the archive files that are going to be transferred to FusionAnalytics. There are also a few buttons on the page as well as a timeline which can be used to check the logs from a different date.

Table Column Headings

Name	Description
Archive File	This is the name of the archive.
Size	This is the size in memory of the archive (MB)
Created	This displays the date and time the archive was created.
Source	The name of the server that the logs are based on.
Archive Comment	This is explanatory text created by the subsystem that created the archive.
Keep	If the text reads Don't Keep then that archive will be removed after a purge has taken place. If it reads Keep then the archive will remain on the drive after a purge has taken place. Simply click the text to change between the two.
FusionAnalytics Connector Lifestyle	Lists where the logs are generated as well as what they were processed by.
FusionAnalytics Connector Status	This will show the status of the connection, No Exception means everything is OK.
FusionAnalytics Connector Actions	Clicking this will transfer the archive from FusionReactor to FusionAnalytics.

Table Buttons

Button	Name	Description
	Show History	Upon clicking this you will be asked if you want to remove the archive, choosing yes will remove it from FusionReactor and the machine that it was stored on.

	Delete Archive	Clicking this will display additional information about the archive. The page is split into two different sections, one contains information about the archive file and the other contains information about how the file was created. Below is two lists summarizing the key points from both sections of the page. **Top section:** • **Archive Filename :** The absolute file path of the archive. • **Original File:** The name of the archive file. • **Timezone:** The timezone used by FusionReactor. • **FusionAnalytics Target:** Specif ies if the archive is to be used by FusionAnalytics if so the FusionAnalytics target is listed. **Bottom Section:** • **Date :** Date the archive was created. • **Source System :** The software used to capture the files for the archive and to create it. • **Subsystem :** The software component within the Source System (FusionReactor) used to perform an action. • **Message :** This is a brief description of the action that the Subsystem performed. |

Rotate Now

To Rotate the logs click on the button that says Rotate Now. This will generate a new archive based on the log information gathered since the previous archive was generated. The archive will then be written to disk.

Clear All Archives

To clear all the archives click on the button that says Clear All Archives. Clicking this will ask you if you want to remove all archives, clicking yes will remove them from the system.

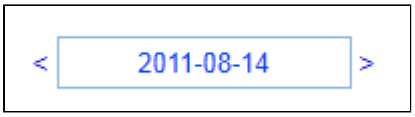
> **WARNING**
>
> Make sure you are 100% sure before removing all archives

Purge Now

To purge the logs click the button that says Purge Now. Selecting this will purge (remove) archives based on rules that have be created by the user. Information about purging can be found in the Log Settings section.

View Logs from a Different Date

Clicking the "<" symbol allows you to revert back to a previous date allowing you to view the details of older logs, clicking the ">" symbol advances the date allowing you to move through time and view more recent logs.



Transfer All

Clicking this will transfer all the archives listed in the table to FusionAnalytics.

FusionAnalytics Settings

This page is used to control how the FusionAnalytics Connector will handle log files for transfer to FRAM or FusionAnalytics DataCollector. The options in the drop down box are listed below in a table.

Configuration	Description
Sent to targets by FRAM	Log files are compressed by FRAM and sent to FusionAnalytics. If FRAM is not available then the option is set to "Sent by this instance" (see below).
Sent to targets by this instance	Log files will be compressed and sent locally, this mode does not require FRAM.
Sent by this instance (uncompressed)	This sends each log file to FRAM or FusionAnalytics directly and uncompressed. After rotation, if this mode is active, the original folder is not deleted. No archives are created.
Disabled	This disables FusionAnalytics Connector meaning no log files will be sent to FusionAnalytics.

After you have selected your option click Save FusionAnalytics Settings.

N.b. Since FRAM itself has no FusionAnalytics targets, the targets to which logs are sent are **always** those of the local instance, regardless of which Connector mode is active.

FusionAnalytics DataCollector Targets

Add FusionAnalytics Targets

Here you can configure your FusionAnalytics targets.

- Adding a New target
- Editing, Verifying and Deleting targets

Adding a new target

Click the Add New Target button and you will be presented with the Add FusionAnalytics Target dialog.

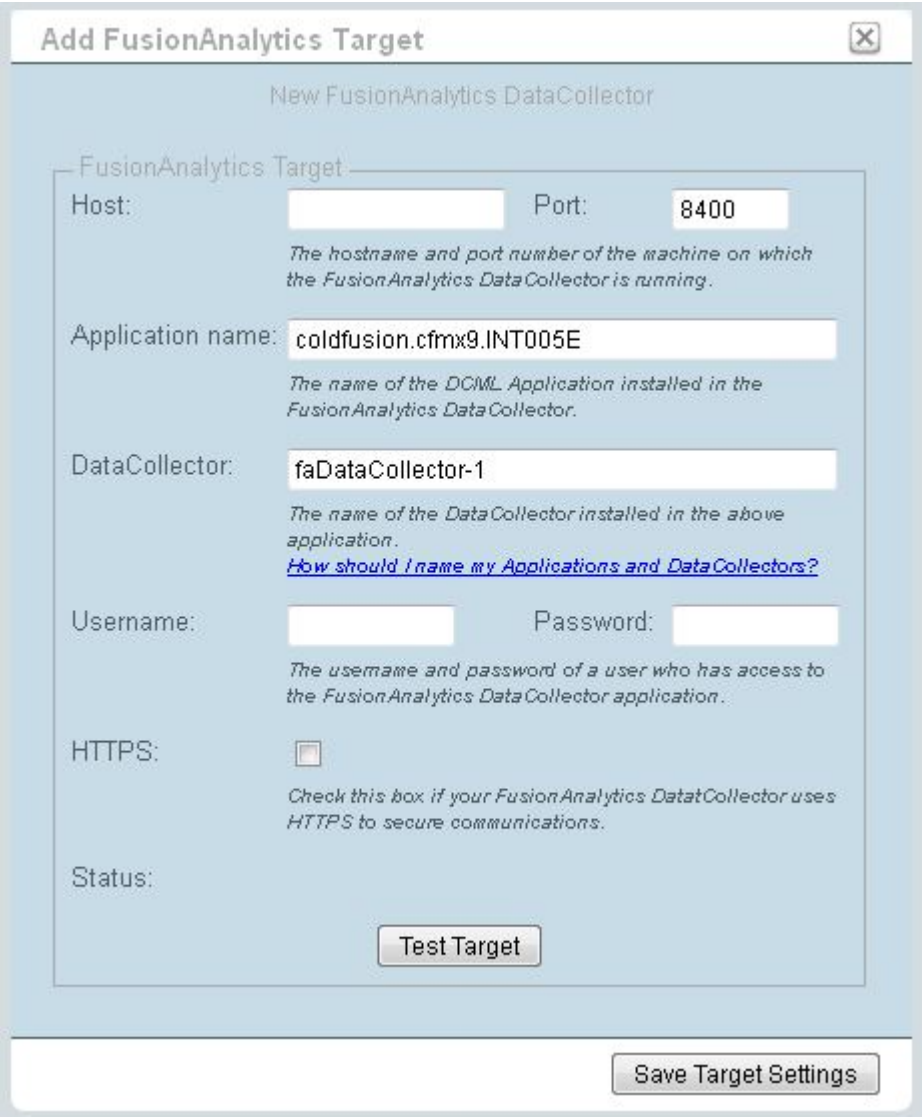


Figure 1: Adding a target

Field	Description
Host	**The machine where FusionAnalytics is installed.** This example uses the local machine, so the IP address is 127.0.0.1 or localhost.
Port	**The machine port.** If you changed the default port during the installation, use that instead; if not, use the default (8400).
Application Name	**The name of the DataCollector Application that you want to collect the data.**
DataCollector	**The name of the actual DataCollector that you want to receive the data.** Each DataCollector application can have multiple DataCollectors that receive data. The default setting is one DataCollector named ***faDataCollector-1***.

Username	**The username for a user who has access to the DataCollector application.** This is only used if the DataCollector requires authentication. If you are using the default setting, this is not needed.
Password	**The password for the above username.**
HTTPS	**Check this box if your DataCollector uses HTTPS.** For more information on this see Setting Up FusionAnalytics with HTTPS.
Status	Once you have entered in the details, click **Test Target** to check that you are connecting to a valid DataCollector.

Click **Test Target**. If the connection was successful the status will display **OK** and you can continue to press Save Target Settings.



Figure 2: Target test successful

Editing, Verifying and Deleting targets

On the FusionAnalytics Targets homepage you will see your targets listed.



Figure 3: Modify targets

Icon	Name	Description
	Edit Target	Allows you to modify the options you entered above.
	Test Target	Tests the connection to the target. The status column will display the results.
	Delete Target	Deletes the target from FusionReactor.

Seeing your data in FusionAnalytics

When you connect FusionReactor to FusionAnalytics, the data will not instantly appear. You will need wait for a transfer to occur.

If you want to force a transfer, in FusionReactor go to **Analytics Connector Status**. Click **Tr ansfer** to send the logs to FusionAnalytics. You will need to wait for each log to be processed by FADC before the data will be available in the Client.

Locked

The FRAM Administrator has the ability to lock various aspect of the individual instance's configuration. This feature was introduced to allow FusionAnalytics settings to be controlled by the FRAM Administrator, rather than the instance administrator, and provides an additional level of access control. With a locked configuration, an instance administrator (or manager) can't change the FusionAnalytics Connector targets and related settings.

Placing the Lock

The lock can **only** be placed in the Instance Manager of FRAM .

Connect to FRAM and open the Instance Manager. Click the *'Edit FusionAnalytics Connector'* button , to open the FusionAnalytics Connector Manager wizard. Hit **Next** to advance to the **Log Settings** page. At the bottom of this page, you'll see a **Lock Settings** section. Check **Locked** to place the lock on the configuration.Make any other changes you need, and advance to the second page. Again, make any changes and click **Save Connector Settings.**

Effects of the Lock

Once placed, some settings - detailed below - are no longer editable from within the FusionReactor Instance. The locked settings will have a lock icon , next to them, and they will be editing of these fields is disabled.

Log Settings

- Log Engine
- Rotation Modee
- Rotate Every/At

FusionAnalytics Settings

- FusionAnalytics Connector Mode

FusionAnalytics Targets

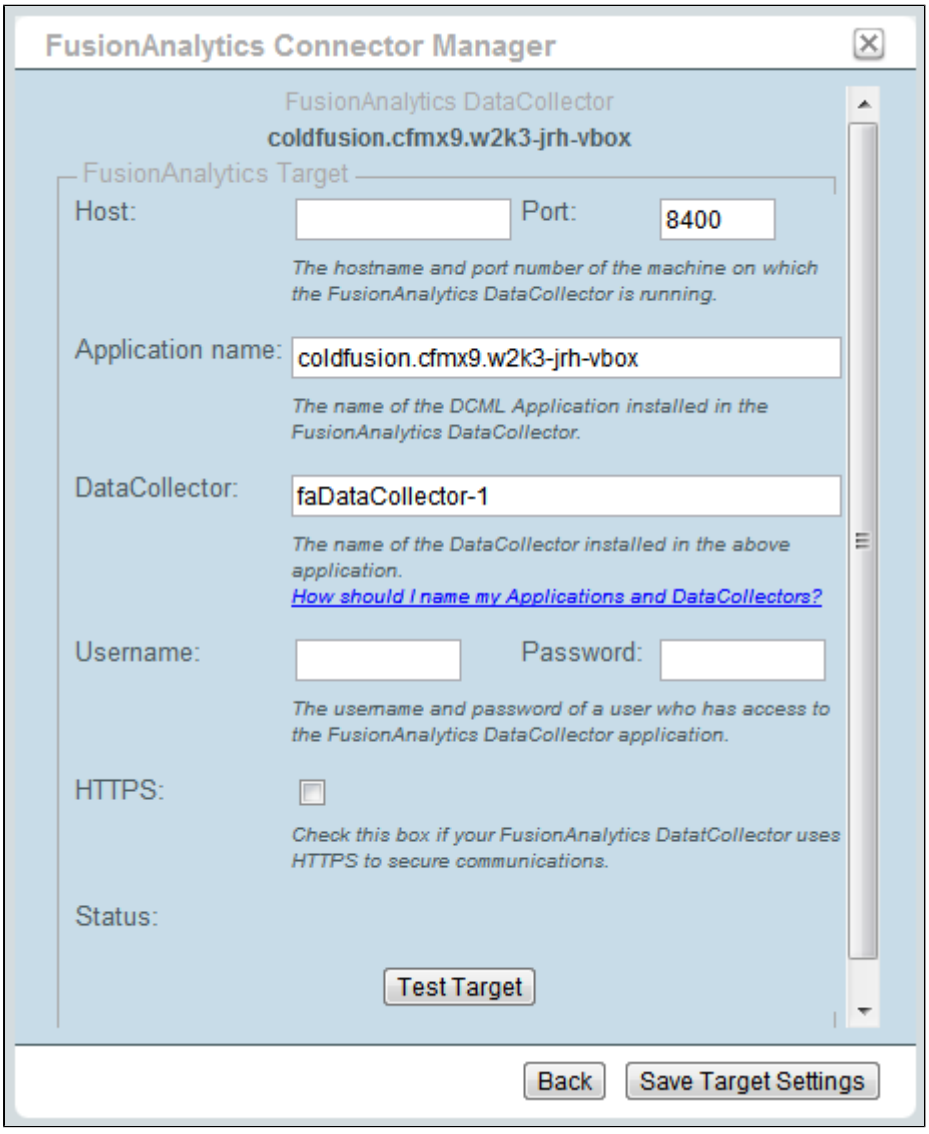
- No targets may be edited, added or deleted.

Naming Applications and DataCollectors

When adding a FusionAnalytics target to FusionReactor's **Analytics -> FusionAnalytics Targets** (normal instance) or within the **FusionAnalytics Connector Manager** (FRAM Instance Manager), you'll be asked for some information about the FusionAnalytics DataCollector target.

An in-depth explanation of FusionAnalytics targets, and the FusionAnalytics DataCollector which defines them, is out of the scope of this article; however it is possible to show you where the data for each field in the wizard comes from.

Here's a capture of a blank form, ready to be filled in.



We'll go through each field and explain where to find it in FusionAnalytics.

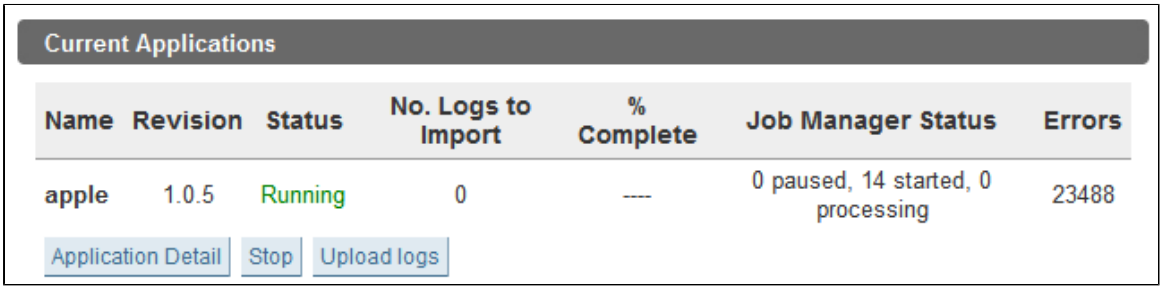
Host and Port

This is the hostname and port number of the FusionAnalytics DataCollector service. If you used the default port, the target will be available on **8400**. The host must be either an IP address or a DNS-resolvable name, which points to a reachable interface, located on the target system.

Application Name

This refers to the name of the individual data collector app running within FusionAnalytics.

You can find it by logging into the **FusionAnalytics DataCollector Administrator,** and looking at the **Applicati on Summary** page. The application name is located in the **Current Applications** table, in the column **Name.** In the figure below, the application **apple** has been defined.



DataCollector

Each **Application** running in FusionAnalytics can define one or more **DataCollectors.** These should be used to cluster together logically-related FusionReactor datasources. For example, all FusionReactor instances on a single machine could feed the same **Application**, but **must** feed different **DataCollectors.**

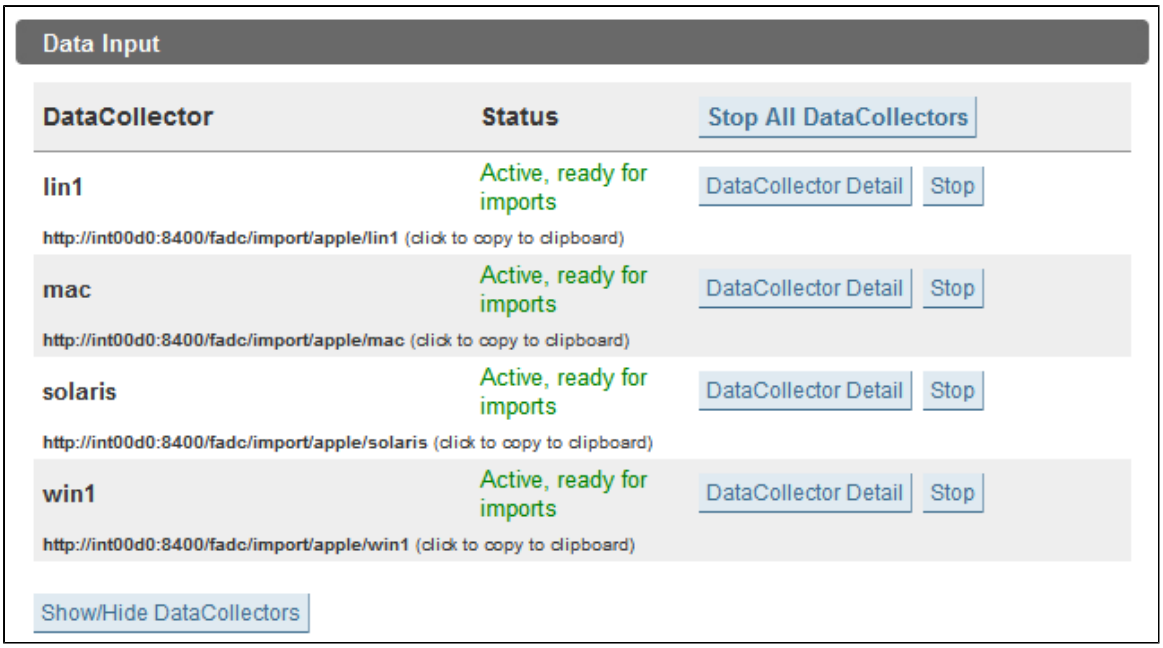
You must not point two different FusionReactor instances at the same **DataCollector**, as this will make the analytics invalid since there is no differentiating field for the different sources.

By using a different **DataCollector,** FusionAnalytics is able to correctly partition the data by including a uniquely-named field (the name of the DataCollector).

If you accepted the default DataCollector name, or you haven't customized your FusionAnalytics DataCollector Application yet, the default DataCollector name will have been set up for you. This is **faDataCollector-1**.

If you're unsure of which names are already in use, you can see which DataCollectors are installed as follows:

Log in to the **FusionAnalytics DataCollector Administrator** and, on the **Application Summary** page click the **Application Detail** button for your application. On the page which appears, click **Show/Hide DataCollectors** to reveal the names of the data collector in use. In the figure below, we've defined four: **lin1, mac, solaris** and **win 1.** .



Pick the one you want to send logfiles to, and enter it's name into the **DataCollector** box in the FusionReactor Connector Manager Wizard.

Username and Password

These are the credentials used to log in automatically to the FusionAnalytics DataCollector. The default username is **admin** and the password is that used to log in to the FusionAnalytics DataCollector Administrator.

HTTPS

Check this box if your FusionAnalytics installation is configured to use **HTTPS** transports.

Test Target

When the box is complete, click **Test Target** to attempt a test connection. If all is well, **OK** appears. If an error occurs, **Error** appears. Clicking on the **Error** opens a window in which diagnostic information can be used to locate the error.

When correct, advance the wizard to the next stage by clicking **Save Target Settings**.

FRAM Target

The FusionReactor Administration Manager (FRAM) installed as part of FusionReactor 4.0.x acts as a vital part in the transfer of logs from individual FusionReactor instances.

FRAM provides a completely isolated instance of FusionReactor which is just for monitoring and instance management, it will continue running even following a JVM crash.

FRAM will also perform process and transfer offloading, making the log transfer to FusionAnalytics a central process performed by a single instance.



Log Transfer & Archiving

By selecting the FusionAnalytics Connector mode **Sent to targets by FRAM** (See FusionAnalytics Settings for more information about log modes) the individual FusionReactor instances the log archiving and transfer to FusionAnalytics will be performed by FRAM. This mode optimizes transfer of archives by having FRAM compress and send files to the FusionAnalytics Targets, offloading the processing power required to compress and transfer of log files onto FRAM. This means that each individual FusionReactor instance will have to perform this costly action, making FusionReactor more transparent on the target server.

The FusionReactor Administration Manager (FRAM) will also act as a central hub for log transfer where multiple instances can send their log files for archiving and transfer via FRAM. By doing this a central archive will be maintained by FRAM containing the log files of all instances using the FusionAnalytics Connector mode **Sent to targets by FRAM** (assuming their FRAM and FusionAnalytics targets are correctly set up. See Edit FusionAnalytics Connector for more detailed information on how to set the FusionAnalytics Connector up).

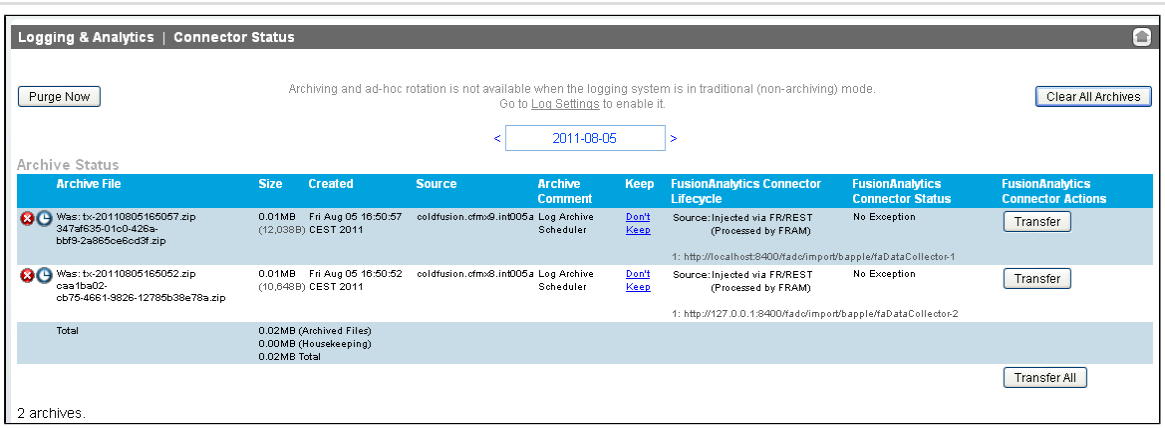


Figure 1: FRAM's Connector Status

Figure 1: FRAM will handle the log transfer and archiving of log files generated by individual FusionReactor instances. Each FusionReactor Instance specifies the FusionAnalytics DataCollector Targets it wishes to send log files to, FRAM will process the log files and send the compressed log files to the specified FusionAnalytics targets.

FRAM will also create a local archive of the log files generated by the instance and save this in its archive folder available at `\instance\FRAM\archive`. By doing this FRAM generates a central storage of files.

More information about each archive can be provided by clicking the *'Archive History'* button next to each listed archive. This will provide information on the life cycle of the archive, all the way from the instance that generated the logs to FusionAnalytics. By having FRAM performing this action the transfer to FusionAnalytics can be easily tracked and maintained by the user of the FusionReactor Administration Manager.

For more information about FRAM please see the online documentation for the FusionReactor Administration Manager (FRAM).

Fusion Reactor

Enabling and Disable

Notifications

Notification Log

Settings

Filter Restrictions

Change Password

Change FR-REST Password

Log

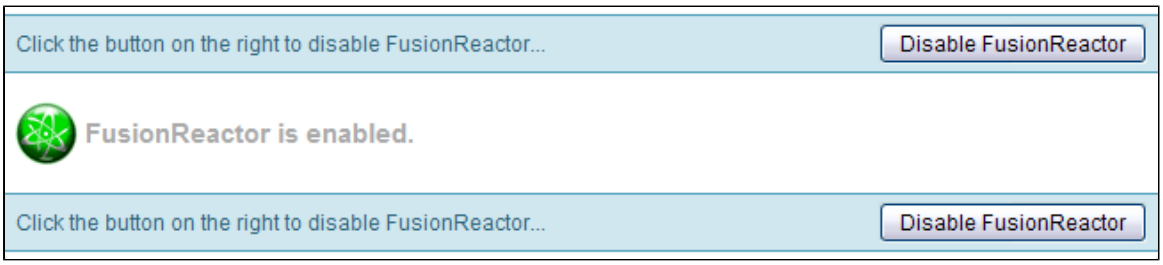
About

Help

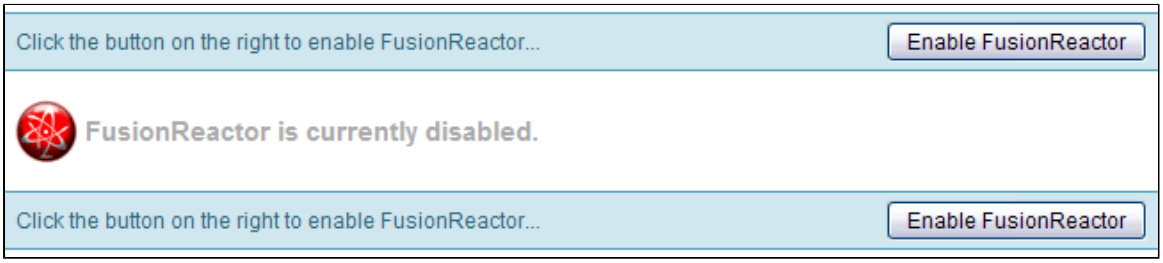
Logout

Enabling and Disable

This page lets you completely enable or disable FusionReactor. Simply click on the "Disable FusionReactor" or "Enable FusionReactor" button to change the state of the system. The current state of the system can be seen from the status icon which is displayed at the bottom of the table of contents. Green means enabled...



...and red means disabled



Notifications

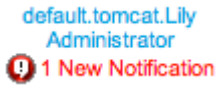
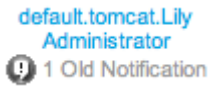
This page allows you to see all the notifications that FusionReactor has raised. Notifications are messages from FusionReactor's various subsystems that alert you to conditions which require attention, or are otherwise abnormal.

You can navigate to this page by clicking on **FusionReactor** in the table of contents then clicking on **Notifications,** or clicking on the text in the table of contents that says **New** (or Old) **Notification** (Figure 1). Notifications are presented in a table (Figure 2) with the following headings:

Name	Description
Origin	The subsystem or plugin which raised the notification.
Time	The time the notification was created.
Title	A succinct title, usually containing a summary of the condition.

Message	A more verbose message. This may contain actions required to relieve the abnormality.
Actions	• **Mark Read** - marks this notification as having been read. It won't appear as a **New Notification** in the table of contents or Enterprise Dashboard (for customers with applicable licenses). The row will be grayed out slightly to indicate this notification has been read. The **Mark Read** button will change to **Mark Unread** - which resets the notification back to its unread state. • **Remove** - removes this notification. Some notifications - usually serious issues - can't be removed. They will be removed by the originating system when the condition to which they pertain clears.

Here are some examples of notifications appearing in the table of contents.

	This notification shows that there is one new (unread) notification requiring attention. The number of old (read) notifications is not displayed in this case. Click on the red text to view the complete list of notifications.
	This notification shows there is one old (read) notification. There are no new (unread) notifications in this case.

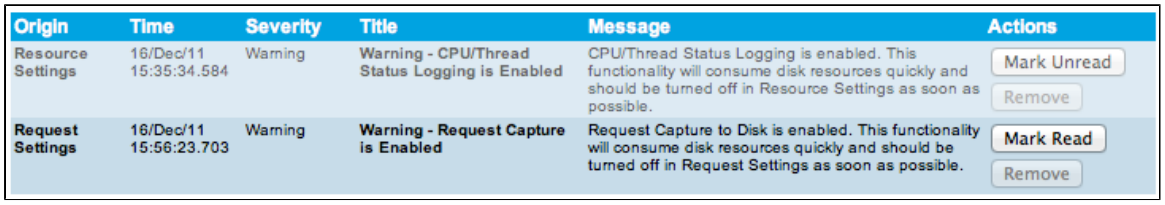


Figure 1: Notification table with one read and one unread notification - both with non-removable notifications.

Global Functions

There are three global functions in the notifications page, represented by three buttons.

> **Changed in 4.0.6**
>
> In FusionReactor 4.0.6, notifications are no longer automatically marked as read when you use the **Notifications** page. You must now explicitly mark them as read using the **Mark Read** or **Mark All Read** buttons.

Mark All Read	Marks all notifications read. They will not appear as **New Notifications** in the TOC indicator, nor will they appear in the Enterprise Dashboard (for customers with applicable licenses).

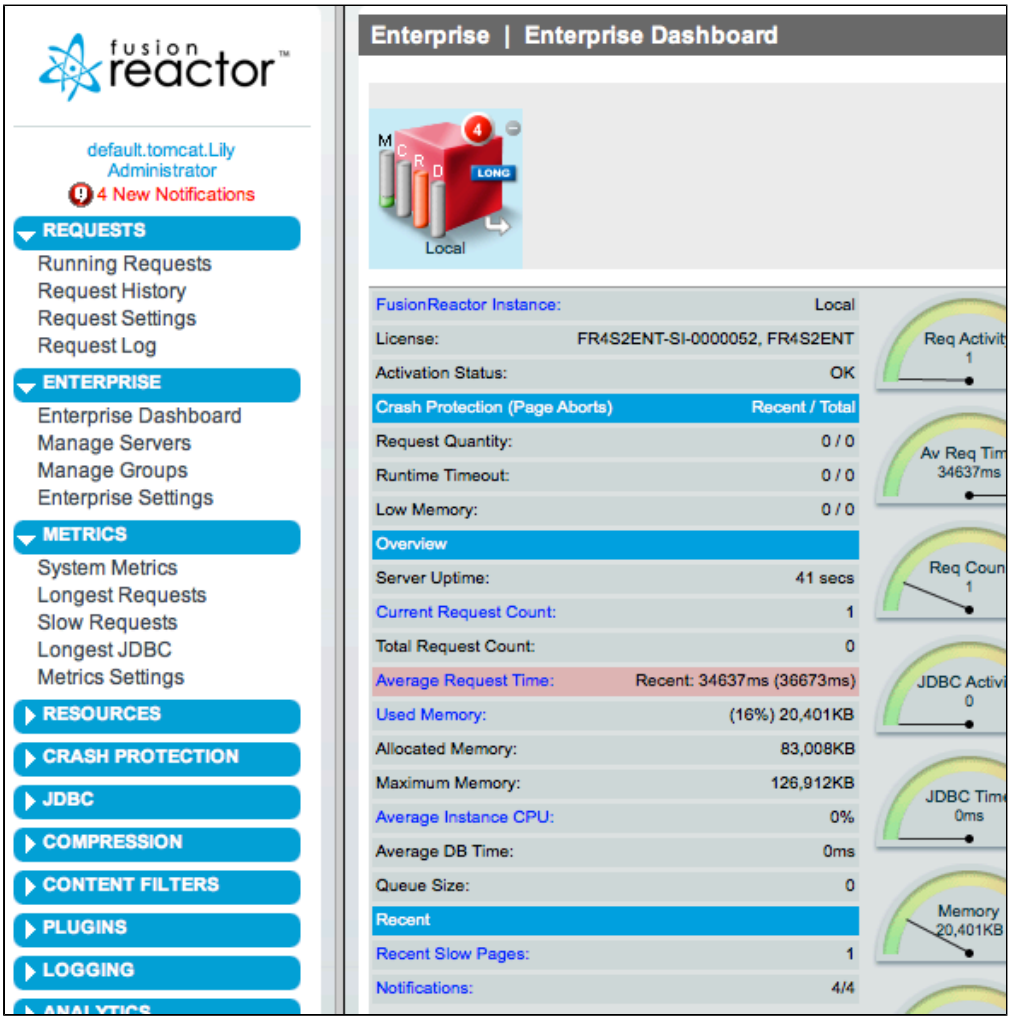
Mark All Unread	Marks all notifications as unread. They will appear as **New Notifications** in the TOC indicator, and will appear in Enterprise Dashboard.
Remove All	Removes all notifications completely.

All buttons will be disabled when they have no logical function; for example if all notifications are unread, **Mark All Unread** will be disabled.

Updating

Since the Notifications list updates dynamically, there will be a short delay (on average 2.5s) until the page updates after a button press. In the meantime, the button text is changed to **"Wait..."** to indicate that FusionReactor has received and processed the request, but that the effect hasn't been reflected in the page yet.

Enterprise Dashboard (New in 4.0.6)



In FusionReactor 4.0.6 Enterprise Edition, notification information is also sent to the Enterprise Dashboard. The cube icon will display a red badge whenever any new notifications are available, and the count of new and total notifications for each monitored server is displayed in the information block.

Notification Log

The notification log is a rotating log file, much like the other logs generated by the system. Fields are space-separated and records are line-separated. Fields which may contain spaces (origin, title and message) are enclosed by double-quotes.

Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date/Time (Milliseconds)	3 [C]	This is the millisecond time at which this request was started or finished (depending upon the Request Status.)
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
Action	6 [F]	A string describing the action associated with this log entry; one of: • **CREATED** - the notification was created and registered with the system. • **MARKED-READ** - the notification was marked as read. • **MARKED-UNREAD** - the notification was marked as unread. • **DELETED-USER** - the notification was deleted by the user • **DELETED-INTERNAL** - the notification was deleted by the system.
Severity	7 [G]	The Severity (text) of the notification
Origin	8 [H]	The name of the creator of this notification.
Title	9 [I]	The title of the notification
Message	10 [J]	The message of the notification
Notification ID	11 [K]	The FusionReactor Notification ID. This is a UUID string assigned to the notification throughout its life.

Exmaple:

Date (Formatted)	Time (Formatted)	Date/Time (Milliseconds)	Version	Server Startup Time (ms)	Action	Severity	Origin	Title	Message	Notification ID

2012-01-25	11:44:20.221	1327488260220	4	1327486508529	CREATED	Info	"frapi-notification.jsp"	"FRAPI-Generated Notification-0"	"This notification was generated by frapi-notification.jsp, a FRAPI test script."	5dc51a0e-32b4-4a63-b101-2fe1a7cf3c87

Settings

The Request Settings page lets you modify email server settings, FusionReactor web root, log file settings, restrictions, stack trace and FR-REST API access.

Email Server

Configuration	Description
From Address	This will be the email address which notification emails will be sent from.
To Address	This is the email address which will receive FusionReactor notifications.
Mail Server	Enter your email server here. This can be entered either as a machine name or an IP address. **NOTE** If your SMTP server does not use port 25 (default), you can specify a port number after the server address/IP address (for example, mail.company.com:587). **NOTE** If your mail server requires authentication, you can specify a user name and password in the format [user]:[password]@[mail server]
CP Email Interval (mins)	This email interval specifies the minimum number of minutes between Crash Protection notification emails. If you are experiencing regular difficulties with a server then use this option to avoid filling your email inbox!

Notification	This lets you completely enable or disabled email notification.

FusionReactor Web Root

Configuration	Description
Web Root	This will be the prefix for FusionReactor requests. If the web root was set to "/fusionreactor/" and you were accessing the internal web server on a local machine over port 8088 then your final URL would look like this: http://127.0.0.1:8088/fusionreactor/

FusionReactor Log File

Configuration	Description
Logging	Choosing Enabled will allow FusionReactor logs to be written, choosing Disabled will prevent them from being written. For information on FusionReactor logs see the Log section.
View Size	This option allows you to limit the number of KB of log that is displayed in the FusionReactor Log. For information on FusionReactor logs see the Log sectio n.
File Size	The size in KB of each of the rotated log files used by FusionReactor.
File Count	The number of rotated log files used by FusionReactor.

NOTE

The File Size and File Count attributes are only available in FRAM

FusionReactor Restrictions

Configuration	Description

Restrictions	This value defines how FusionReactor Restrictions behave. • **Monitor requests that match the rules** - By default, no requests will appear in FusionReactor. If a request matches one of the rules defined on the Filter Restrictions page then it will appear. • **Ignore requests that match the rules** - By default, all requests will appear in FusionReactor. If a request matches one of the rules defined on the Filter Restrictions page then it will be invisible.

Stack Trace : (Configuration Item only appears when using a Java 1.5 or higher JVM)

Configuration	Description
Stack Trace Method	Lets you select the method used by FusionReactor to generate and capture stack traces. Older versions of Java (before Java version 1.5) could not easily generate and capture stack traces of the running JVM but FusionReactor adds this functionality even on to such older versions. It is recommended to use the **Java's Stack Trace** method on Java version 1.5 or higher. On older versions of Java this configuration item does not appear because FusionReactor's default stack trace method is the only method that can be used to generate stack traces. • **Use Java's Stack Trace** - Use Java's built in implementation of generating stack traces • **FusionReactor Default** - Use FusionReactor's default stack trace method

FR-REST API Access

This is a HTTP based API that is used to communicate commands and retrieve data from FusionReactor. This API allows FusionAnalytics Connector system to send logs to FusionReactor Administration (FRAM).

Configuration	Description

Enable FR-REST API	• **Enable locally, no authentication :** This enables the API to work on a local machine with no authentication. For information on creating or changing a password see the Change FR-REST Password section. • **Enable locally with authentication :** This enables the API to work on a local machine with authentication. This option won't be presented until you create a FR-REST password, for information on how to do so see the Change FR-REST Password section. • **Enable globally, with authentication :** This enables the API to work on a global machine (other machine in the network) with authentication. This option won't be presented until you create a FR-REST password, for information on how to do so see the Change FR-REST Password section. • **Disable :** This disables the FR-REST API from working. **NOTE** In order for FusionAnalytics to be able to work alongside FusionReactor, this option can **NOT** be disabled.
FR-REST User	You can input the name of the FR-REST User here, you can change it as you please. It is also recommended that you also set a password, for help doing this see the Change FR-REST Password secti on.

Filter Restrictions

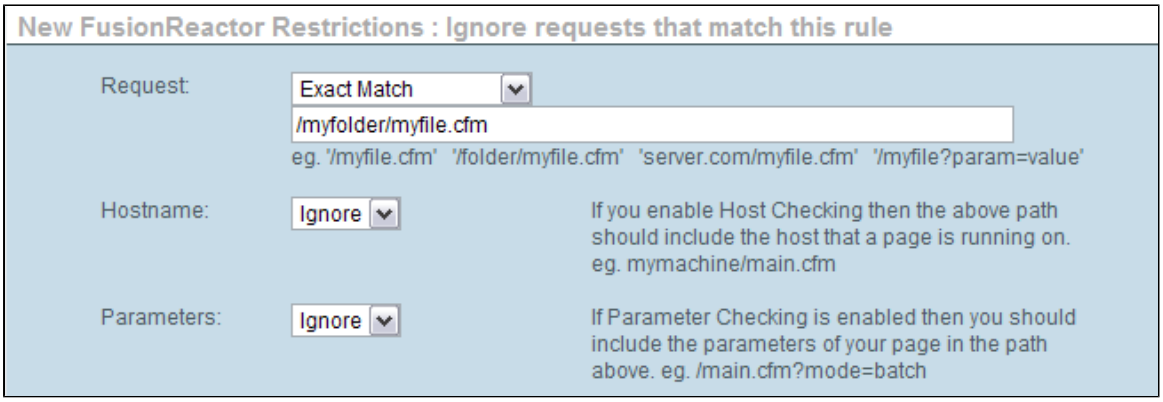
FusionReactor restrictions are set up so that FusionReactor can ignore certain requests, or alternatively, you can specify a set of pages which you want to monitor with FusionReactor while all others are ignored. To select which mode the Restrictions page works in, change the Restrictions option on the FusionReactor Settings page. The FusionReactor Restrictions page has two main areas: At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for the application.

New FusionReactor Restrictions

To enter a restriction you will need to fill the following fields:

Configuration	Description

Request	This is the page you want to exclude. Firstly select if you are creating an "Exact Match" or "Regular Expression". **Exact Match** Next enter the location of your page. By default the request will be an exact match and should look something like "/batchStats.cfm" or "/batch/dailyStats.cfm" (beginning with a '/' and without parameters.) . The text you input here will depend on the options you choose for the Hostname and Parameters fields. **Regular Expressions** A good use of Regular Expressions here would be to restrict everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Content Filters. The text you input here will depend on the options you choose for the Hostname and Parameters fields.
Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you enable Hostname checking then the request field should also include a hostname and look something like "productionMachine/batchStats.cfm". You should NOT include HTTP or HTTPS at the start of this string.
Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the Hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example "/batchStats.cfm?period=YEAR". Once again, Regular Expressions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.



Once you submit the form the page will be refreshed and your new restriction will be listed at the bottom of the

page.

Managing Existing Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

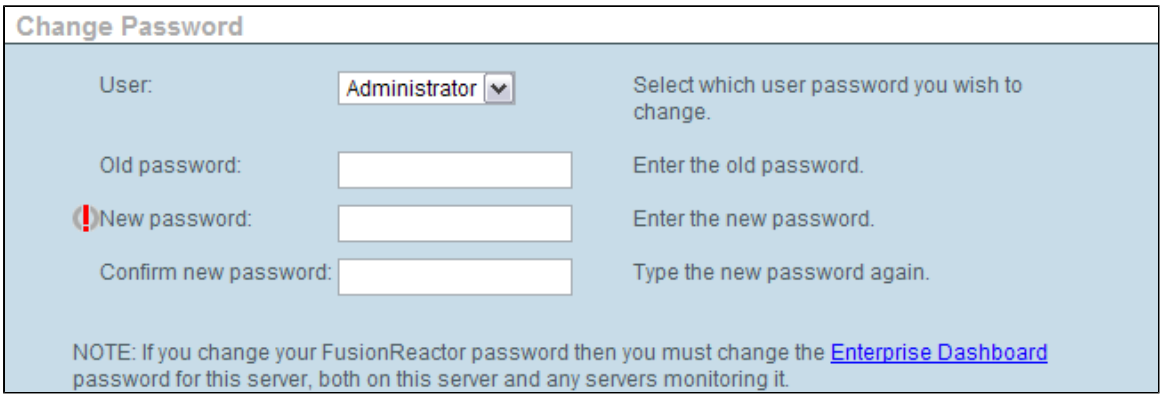
Button	Name	Description
	Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above.
	Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.



Change Password

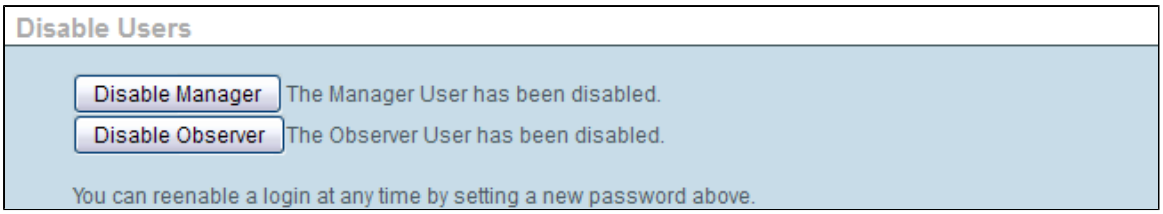
Change Password

From this page you can change the password for the various users within the system. To change a password, select the user type and enter a new password (plus verification.)



Disable a User Account

If you want to disable a user type then click on one of the "Disable [user type]" buttons.



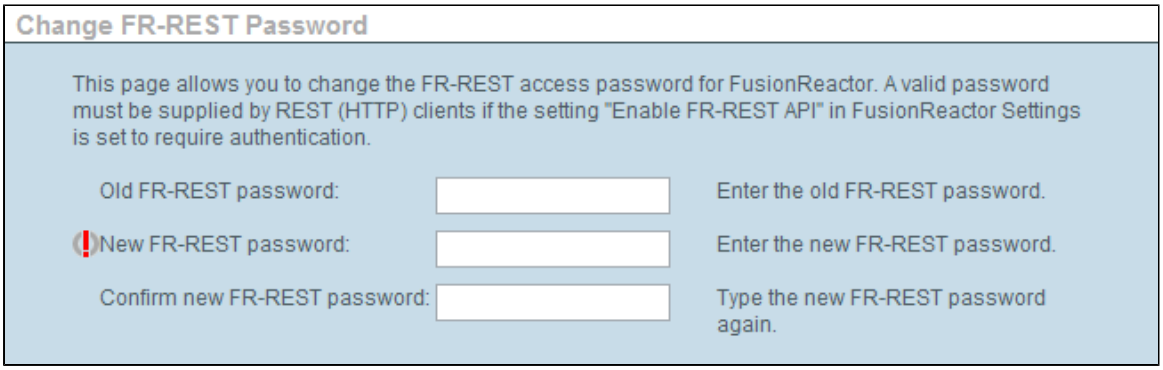
Enable a User Account

To enable a user type which has been disabled, just set a new password for that user type.

Change FR-REST Password

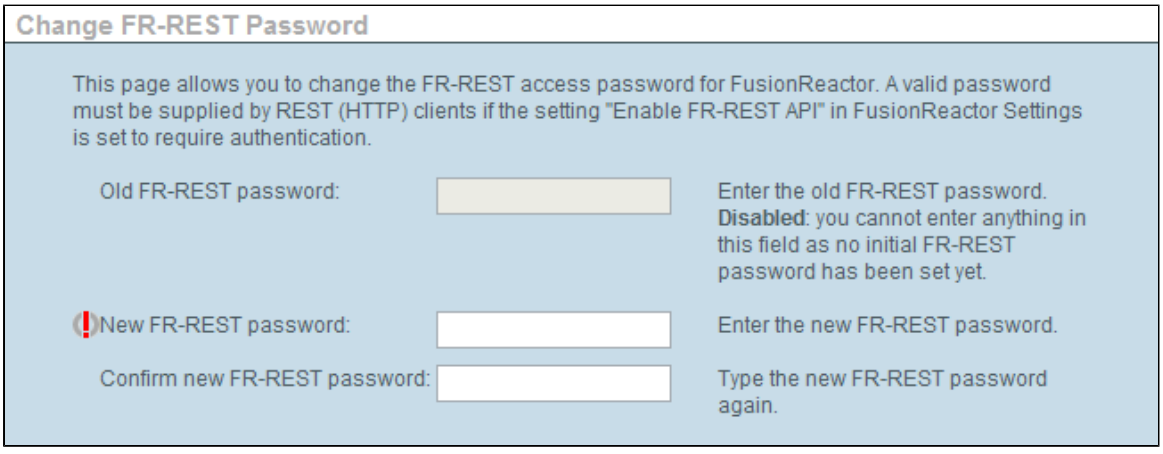
Change FR-REST Password

To change the FR-REST password simply type in the old password (Old FR-REST password field), then type in a new one below (New FR-REST Password) and verify it below (Confirm new FR-REST password) and click the button that says Save FR-REST Password.



Create FR-REST Password

If you have never created a FR-REST password you will not be able to type anything into the Old FR-REST Password field, just proceed to filling in both the text fields below and click the button that says Save FR-REST Password.



See Also

Settings - FR-REST API

Log

The FusionReactor Log page is a simple display page which shows the contents of the FusionReactor Log file. The file contains information about FusionReactors configurations and activity. The FusionReactor Logs are a

set of rotating files which you can configure from the FusionReactor Settings page. By default these files are located in the /FusionReactor/instance/[your_instance_name]/log directory.

Deleting the FusionReactor Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the FusionReactor Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

About

In the About page in FRAM you can do the following:

- Upload a license
- Activate a license
- Activate a license offline
- Check for updates by clicking on the link titled "check For FusionReactor updates", this takes you to a webpage that tells you if your version is up to date. If it is outdated you will be given information on how to upgrade.



About page in FusionReactor |

Types of licenses

For information on the different kinds of licenses you can purchase check out the comparison page on the website. For information on how to purchase licenses check out the FusionReactor website.

Activation Period

Once you upload your license you have 10 days to activate it, for information on how to do so see the Activation Period section.

Activating Multiple-Key License Files

One license file can contain many keys meaning you don't need several license files if you intend to use FusionReactor on many machines, for more information on this see the Activating Multi-Key License Files sectio n and the License Policy section.

Also See

Uploading your License File
Activation Period

License Activation
Activating Multi-Key License Files
License Policy
Install FusionReactor Instance
License Purchase Page

Help

Clicking on this link will take you to the help system which you are reading right now!

Logout

Click the "Logout" link to exit the FusionReactor Administrator and be returned to the login page.

Request Detail

Stack Tracing

Request Detail Page

Stack Tracing

Stack traces are split into two sections. Firstly comes a set of information associated with this thread, or request. Please note that you will see more information here when Stack Tracing a request rather than a thread.

After these information fields you will see the Stack Trace itself. Note that the format of Stack Traces may alter depending on your operating system or Java version. Generally you will see a list of class names, method calls and line numbers. The most current list item will be listed first.

Killing a Thread or Request

If you click on the "Kill" button then you will see a prompt asking if you are sure. Click the "OK" button in this dialog box to kill the request or thread.

> **WARNING**
>
> The "Kill" button does not give a request the opportunity to complete in it's own time, it forces the request to end immediately. Using this feature can cause the JVM to become unstable. Also, while FusionReactor does allows you to stop any thread, it is highly inadvisable to stop any threads in the 'system' thread group, or any threads integral to the running of your J2EE container.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Request Detail Page

You can get to the Request Detail page from any page which lists requests: Running Requests, Request History, Longest Requests, Slow Requests, and pages linked from Request Metrics. Clicking the Request Details logo ( ) by a request will take you to a page which contains the following tabs:



The Tabs

At the top of the Request Details page you will see a copy of the request line from the page you came from. Underneath that there are six tabs.

Name	Description
Main	The Main tab shows general information about a request such as script name, execution time, memory levels and flags.
Headers	The Headers tab shows the HTTP headers for the request and response objects.
Cookies	The Cookies tab takes the Cookie string from the Headers tab and renders it in a more readable format along with other cookie information such as maximum age and whether is it secure.
JDBC	The JDBC tab will show any queries which have been run by your page. If you also have User Trace Markers on this page then you can click the "Include Markers" link to create a merged view of both markers and queries. At the bottom of the JDBC tab you will find the total and average times for queries running on this page.
Markers	The Markers tab shows all User Trace Markers which were called during this request along with the amount of time which elapsed between them.
AMF	If Flash Remoting Support is enabled (Request Settings) then the AMF tab displays BOTH the REQUEST and the RESPONSE which came in AMF format. This is an essential tool for debugging RIA's. AMF is typically used in Flash Remoting and in applications such as Flex. For further information, refer to the AMF Dissector Plugin

Refreshing the Page

You can refresh the page at any time by pressing the refresh button which is located at the top right corner of the Request Detail page.

Also See

AMF Dissector Plugin

Analyzing Data

FusionAnalytics

Importing and Graphing Data in Excel

FusionAnalytics

FusionAnalytics

The best and most efficient way to analyze FusionReactor data is to use FusionAnalytics.

FusionAnalytics is a ColdFusion Application and server analysis tool, which integrates seamlessly with FusionReactor. FusionAnalytics is all about "Making IT Better" and giving you the information and data to help you make better business decisions, improve application performance and quality of your applications as well as measure exactly how your applications are performing on a continuous day to day basis.

For more information about analyzing FusionReactor data using FusionAnalytics see the Analytics menu item providing the current status between your instance of FusionReactor and FusionAnalytics.

Importing and Graphing Data in Excel

Most FusionReactor Log files can be easily imported into Excel (and other tools that can import delimited files) in order to analyze them or create graphs.

There are 4 major log files that you can import:

- request-*.log
- resource-*.log
- jdbc-*.log
- crashprotection-*.log

If you are using the FusionReactor Extensions for ColdFusion (FREC) plugin then you will also have access to the following log files:

- dbpoolstats.log
- hitcountstats.log
- plugin-fr-logrotator-plugin.log
- realtimestats.log

> **NOTE**
>
> All of the log file information can be imported and analyzed directly into FusionAnalytics

How Do I import a FusionReactor Log File into Excel

1. Make a copy of the log file you want to import. (It's better to work on copies than the live log files otherwise you may stop them from being written to)
2. Start Excel and from the [File] menu select [Open...] and then the log file. You may have to select Files of type: All files (*.*) to see the log files in the file picker
3. The Text Import Wizard dialog will appear
4. For the Original data type choose : Delimited and then press [Next >] at the bottom of the dialog
5. Select Space as the Delimiter and then press [Finish] at the bottom of the dialog
6. The Spreadsheet should now be filled with data in separate columns
7. Click on top of one or more columns that you wish to graph (to highlight the entire column) and then click on the small graph icon (Chart Wizard) in the toolbar (or from the menu [Insert]->[Chart]). Note you can find out what data the Excel columns contain by using the field number [letter] in the descriptions of the log files above.
8. Choose the graph style (e.g. line graph, scatter graph or bar graph) and then click finish.

> **NOTE**
>
> You can find out what data is stored in the columns using the descriptions of the log files above or by looking into the request-headers.txt, resource-headers.txt, jdbc-headers.txt or crashprotection-headers.txt files stored in the folder with the log files.

Users

[Types Of User Account]

[Managing The User Accounts]

Types Of User Account

FusionReactor supports three different types (or levels) of user account:

- Administrators
- Managers
- Observers

> **NOTE**
>
> FusionReactor does not have individual users with user names and passwords, but three fixed accounts each of which has its own password.

Administrators

Administrator is the most powerful type of user account in FusionReactor. They are able to access all of the functions available in FusionReactor: (Note: some features are only available on Enterprise Versions)

- Creating and Maintaining new instances of FusionReactor on a computer
- Configuring the FusionReactor instances visible to the Enterprise Dashboard (Enterprise only)
- Configure Crash Protection and all other settings for an instance of FusionReactor
- Managing All User Accounts
- Uploading and Activation of License Keys
- Examine All Request and Metric Data
- Kill Running Requests

Who should use the Administrator Account?

The Administrator should be used by users who need to create and maintain the instances of FusionReactor on computer, setup and manage the instances that are visible on any Enterprise Dashboard and maintain the user accounts. For example, in hosting environments this account might be used by the hosting company to configure FusionReactor for customers on their application server instances.

Managers

Managers are "Instance Administrators". They Manage or Administer a single instance of FusionReactor. They are able to: (Note: some features are only available on Enterprise Versions)

- Configuring the FusionReactor instances visible on this instance's Enterprise Dashboard
- Add instances visible on this instance's Enterprise Dashboard (if they have the Administrator account password for the instance they wish to add)
- Online and Offline Servers in the Enterprise Dashboard
- Configure Crash Protection and all other settings for this instance of FusionReactor
- Managing Their User Account and the Observer Account
- Examine All Request and Metric Data for this instance
- Kill Running Requests

Who should use the Manager Account?

The Manager account should be used by users who need to manage or administer an application server. For example, in hosting environments this account might be used by the customer's application administrator.

Observers

Observers are only able to view the following information on a specific instance: (Note: some features are only

available on Enterprise Versions)

- View the Running Request
- View the Request History
- View the Request Log
- View the Enterprise Dashboard
- View the System Metrics
- View the Longest Requests, Slow Request and Longest JDBC Request Reports
- View the Resource Reports and Graphs
- View the Resource Log, Crash Protection Log and JDBC Log
- View the Help and About Information

Who should use the Observer Account?

The Observer account should be used by users who need to examine the performance and metrics from an application server. For example, in hosting environments this account might be used by the customer's application developers to examine how well the application performs in production.

Enterprise Dashboard

When a user clicks on a link on the Enterprise Dashboard for a monitored server they are transferred to that server instance of FusionReactor with the same level of user account that they are logged in as. For example, this means that a Manager on an instance that is monitoring another instance of FusionReactor will be granted Manager level access to the other instance of FusionReactor.

Managing The User Accounts

FusionReactor has three user accounts (Administrator, Manager and Observer).

Enabling a User Account

The user accounts are enabled by entering a password for the account.

> **NOTE**
>
> The Administrator account is always enabled and must have a password assigned. Note that when you enter a password and click on the [Save Password] button the account becomes active immediately.

Disabling a User Account

The Manager and Observer accounts can be disabled using the [Disable Manager] and [Disable Observer] buttons the Change Password screen. An account will be disabled immediately after clicking on the Disable button for the account. To enable an account again you must enter a password for the account.

Enabling Accounts (Setting Passwords) During Installation

During the installation process you must enter a password for the Administrator account, but you can also enter a password for the Manager and Observer accounts. If you do not enter a password for an account during the installation process, the account will be disabled. You can enable the account at any moment by entering a password for the account on the Change Password screen.

NOTE

When you enter password and click on the [Save Password] button, the account becomes active immediately.

Regular Expressions

Regular Expressions and Restrictions

Regular Expressions and Content Filters

Regular Expressions and Restrictions

Regular Expression Restrictions for FusionReactor use the Java 1.4 RegEx language. For more technical information on Regular Expressions please refer to the online Java documentation

Quick Overview of RegEx Special Characters

RegEx	Description
.	matches any character
\d	matches any digit
\w	matches any word character [a-zA-Z_0-9]
\s	matches any whitespace character
[]	contents define a character class
[abc]	matches a, b or c
[^abc]	matches anything except a, b or c
|	alternation operator ('or')
X|Y	matches X or Y
X?	matches X once, or not at all
X*	matches X zero or more times (matching as much as possible)
X*?	matches X zero or more times (matching as little as possible)
X+	matches X one or more times
()	grouping characters

^	start of line
$	end of line
(?i)	turn on case insensitive matching

Example Restrictions

RegEx	Description
/myfolder/myfile.cfm	This will match any URL which contains "/myfolder/myfile.cfm"
/myfolder/(.*)	This will match anything within "myfolder"
/myfile\.jsp	This will match any files named myfile.jsp which are inside a directory.
\.((cfm)|(jsp))	This will match any .cfm or .jsp file.
/myfile\.cfm/?(.*)mode=add	This expects parameter checking to be enabled. It will match any file named myfile.cfm which has the string "mode=add" as part of it's parameters. (Note: This will also match mymode=add.)
/batch[0-9].cfm	This will match files names batch0.cfm, batch1.cfm ... batch9.cfm
(?i)/batch[0-9].cfm	This will match files names Batch0.cfm, BATCH1.cfm ... batch9.cfm

Common Issues when using with FusionReactor

When using Regex Expressions with FusionReactor there are a few common problems.

FusionReactor allows you to perform matches with or without the Hostname by setting the Hostname field to Check or Ignore respectively. When the Hostname field is set to Check you should include the hostname but not the protocol (e.g. http) or the leading (://). If you set the Hostname flag to Ignore then you typically need to start the Regex Expression with a / to match the whole URL from the root. Generally it is recommended to set the Hostname to Ignore and remove the hostname and protocol from the Regex because request could be made against IP addresses as well as hostnames.

RegEx	Hostname	Description
my.host.com/myfolder/myfile.jsp	Checked	This will match the file /myfolder/myfile.jsp when the request was made against the host my.host.com
/myfolder/myfile.jsp	Ignore	This will match any URL /myfolder/myfile.jsp made against any host or IP address

Often URL's are not case sensitive so it's generally a very good idea to set the Regex case insensitive flag (?i) on your Regex expressions.

RegEx	Hostname	Description
(?i)/myfolder/myfile.jsp	Ignore	This will case insensitive match for any URL /myfolder/myfile.jsp made against any host or IP address

Regular Expressions and Content Filters

Regular Expression Restrictions for FusionReactor use the Java 1.4 RegEx language. For more technical information on Regular Expressions please refer to the online Java documentation.

> **NOTE**
>
> Since the regular expression buffer operates continuously, it has no concept of line breaks; the '$' (end of line) and '^' (start of line) match characters are not available. Also, pattern grouping and the backref operator '\n' are not available.

Quick Overview of RegEx Special Characters

RegEx	Description
.	matches any character
\d	matches any digit
\w	matches any word character [a-zA-Z_0-9]
\s	matches any whitespace character
[]	contents define a character class
[abc]	matches a, b or c
[^abc]	matches anything except a, b or c
|	alternation operator ('or')
X|Y	matches X or Y
X?	matches X once, or not at all
X*	matches X zero or more times (matching as much as possible)

X*?	matches X zero or more times (matching as little as possible)
X+	matches X one or more times
()	grouping characters

Example Search and Replace Filters

RegEx	Description
FusionReactor	This will match the exact string "FusionReactor" anywhere within a page.
[Ff]usion[Re]eactor	This matches FusionReactor without needing to capitalize the F and the R
[Ff][Uu][Ss][Ii][Oo][Nn]	This matches the word "Fusion" in any case.
<a href="http://www.machine.com(.*?)">	This will match any link to a specific machine (which you could replace with a different link to a "Not currently available" page or something.
<meta http-equiv="refresh"(.*?)>	This will match meta refresh tags which you could then remove.

FusionReactor API (FRAPI)

What is FRAPI?

FRAPI is the FusionReactor API and it gives you the ability to access FusionReactor functionality from directly within your ColdFusion pages.

What can I do with FRAPI?

With FRAPI you can make your pages aware of what else is currently running on your server. You can change FusionReactor settings and you can set user trace markers.

What is a User Trace Marker?

A user trace marker can be used for many purposes. For example, you could use it to debug a page and work out where the flow of execution goes, or you could surround critical sections of your code and work out exactly how long they are taking to execute. The main advantage of this mechanism is that you can implement it without having to display any debug information to end users or having to write to hard-to-read debug files on your server.

FRAPI Examples in ColdFusion

FRAPI Examples in ColdFusion

The FRAPI lets you call FusionReactor methods directly from your ColdFusion pages. JavaDocs are available for the FRAPI and are available from the FusionReactor site. Click here to view the FRAPI JavaDocs.

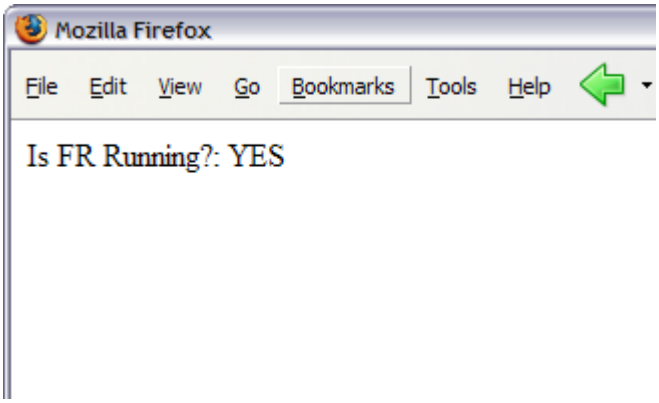
Lets first start with some simple examples though...

This is one of the simplest FRAPI calls:

```
<cfset frapiClass = createObject("java",
"com.intergral.fusionreactor.api.FRAPI")>
<cfset frapi = frapiClass.getInstance()>

<cfoutput>Is FR Running?:
#frapi.isFusionReactorRunning()#</cfoutput
>
```

This page simply outputs a YES/NO value depending on the running state of FusionReactor.



Using FRAPI to look at running requests

Now for something a little more interesting:

```
<cfset frapiClass = createObject("java",
"com.intergral.fusionreactor.api.FRAPI")>
<cfset frapi = frapiClass.getInstance()>
<cfset dateClass = createObject("java",
"java.util.Date")>

<!--- Loop over each running request --->

<cfset requestArray =
frapi.getRunningRequests()/>
<hr>
```

```
<cfloop index="i" from=1
to=#ArrayLen(requestArray)#>

   <cfset dateObject =
dateClass.init(requestArray[i].getStartTim
eMillis())>
   <cfoutput>
      ID:
#requestArray[i].getFusionRequestId()#<br/
>
      URL:
#requestArray[i].getRequestUrl()#<br/>
      Start: #dateObject.toString()#<br/>
      <hr>
   </cfoutput>

</cfloop>

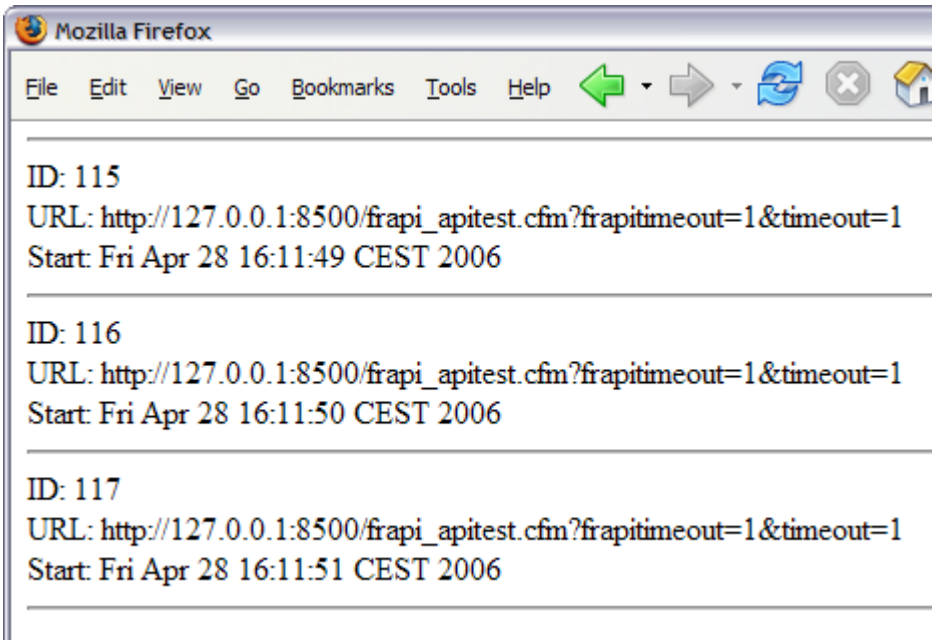
<!--- Sleep for 5 seconds --->

<cfobject type="java" action="create"
class="java.lang.Thread" name="Thread">

<cfscript>
```

```
    Thread.sleep( JavaCast( "long", 5000 )
);
</cfscript>
```

This bit of code will print out a list of all requests currently running. (We've added a wait at the bottom so we can run it several times and get a result something like this:



The FusionRequestSurrogate

Calling frapi.getRunningRequests()returns an array of FusionRequestSurrogates. These objects contain proxied information from an internal FusionRequest (web request). It is safe to store references to this object, since the data it encapsulates is copied from the underlying FusionRequest. There is a lot of data that you can get about a request:

```
<cfscript>
    frapiClass = createObject("java",
"com.intergral.fusionreactor.api.FRAPI");
    frapi = frapiClass.getInstance();
</cfscript>

<table border="1">

<tr>
   <td title="getBytes()"
>Bytes</td>
```

```
   <td title="getCfid()"
>Cfid</td>
   <td title="getCftoken()"
>Cftoken</td>
   <td title="getCpuUsage()"
>CPU</td>
   <td title="getCrashProtection()"
>CP</td>
   <td title="getCrashProtectionStatus()"
>CPStat</td>
   <td title="getExecutionTime()"
>ETime</td>
   <td title="getFusionRequestId()"
>FRid</td>
   <td title="getIpAddress()"
>IP</td>
   <td title="getJavaSessionId()"
>JSid</td>
   <td title="getMethod()"
>Met</td>
   <td title="getOutputCompressorState()"
>OCS</td>
   <td title="getQueryString()"
>QS</td>
   <td title="getRequestId()"
>Rid</td>
   <td title="getRequestStatus()"
>RS</td>
   <td title="getRequestUrl()"
>RURL</td>
   <td title="getRuntime()"
>RTIME</td>
   <td title="getStartTimeMillis()"
```

```
>STM</td>
   <td title="getStatusCode()"
>SC</td>
   <td title="getThread()"
>Tid</td>
   <td title="getThreadCpuTime()"
>TCPU</td>
   <td title="getTimeClosed()"
>CLOSE</td>
   <td title="getTimeOfDeath()"
>Dead</td>
   <td title="getTimeOpened()"
>OPEN</td>
   <td title="getTtfb()"
>Ttfb</td>
   <td title="getTtlb()"
>Ttlb</td>
   <td title="isImmortal()"
>I</td>
   <td title="jdbcGetRunningStatement()"
>JRS</td>
   <td
title="jdbcGetRunningStatementDatasource()
"     >JDS</td>
   <td
title="jdbcGetRunningStatementStartTimeMil
lis()">JRST</td>
   <td title="jdbcTotalExecutionTime()"
>TET</td>
   <td title="jdbcTotalQueryCount()"
>TQC</td>
   <td title="jdbcTotalRows()"
>TR</td>
```

```
   <td title="jdbcTotalTime()"
>RR</td>
   <td title="wasQueued()"
>Q</td>
</tr>

<cfset requestArray =
frapi.getRunningRequests()/>

<cfloop
   index = "i"
   from = 1
   to = #ArrayLen(requestArray)#>

   <cfoutput>
   <tr>

<td>#requestArray[i].getBytes()#</td>
      <td>#requestArray[i].getCfid()#</td>

<td>#requestArray[i].getCftoken()#</td>

<td>#requestArray[i].getCrashProtection()#
</td>

<td>#requestArray[i].getCrashProtectionSta
tus()#</td>

<td>#requestArray[i].getExecutionTime()#</
td>

<td>#requestArray[i].getFusionRequestId()#
</td>
```

```
<td>#requestArray[i].getIpAddress()#</td>

<td>#requestArray[i].getJavaSessionId()#</
td>

<td>#requestArray[i].getMethod()#</td>

<td>#requestArray[i].getOutputCompressorSt
ate()#</td>

<td>#requestArray[i].getQueryString()#</td
>

<td>#requestArray[i].getRequestId()#</td>

<td>#requestArray[i].getRequestStatus()#</
td>

<td>#requestArray[i].getRequestUrl()#</td>

<td>#requestArray[i].getRuntime()#</td>

<td>#requestArray[i].getStartTimeMillis()#
</td>

<td>#requestArray[i].getStatusCode()#</td>

<td>#requestArray[i].getThread()#</td>

<td>#requestArray[i].getThreadCpuTime()#</
td>
```

```
<td>#requestArray[i].getTimeClosed()#</td>

<td>#requestArray[i].getTimeOfDeath()#</td
>

<td>#requestArray[i].getTimeOpened()#</td>
      <td>#requestArray[i].getTtfb()#</td>
      <td>#requestArray[i].getTtlb()#</td>

<td>#requestArray[i].isImmortal()#</td>

<td>#requestArray[i].jdbcGetRunningStateme
nt()#</td>

<td>#requestArray[i].jdbcGetRunningStateme
ntDatasource()#</td>

<td>#requestArray[i].jdbcGetRunningStateme
ntStartTimeMillis()#</td>

<td>#requestArray[i].jdbcTotalExecutionTim
e()#</td>

<td>#requestArray[i].jdbcTotalQueryCount()
#</td>

<td>#requestArray[i].jdbcTotalRows()#</td>

<td>#requestArray[i].jdbcTotalTime()#</td>

<td>#requestArray[i].wasQueued()#</td>
   </tr>
   </cfoutput>
```

```
</cfloop>

</table>
```

Setting a request's crash protection timeout

The following code will change the timeout setting for Crash Protection for this request:

```
<cfset frapiClass = createObject("java",
"com.intergral.fusionreactor.api.FRAPI")>
<cfset frapi = frapiClass.getInstance()>

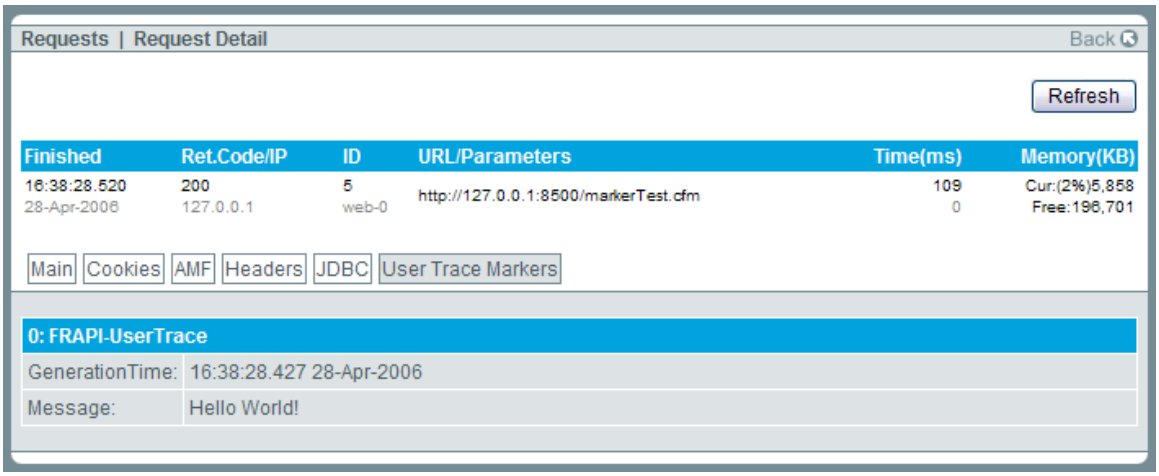
<!--- Force Crash Protection to 10 seconds
for this page --->
<cfset frapi.setPageTimeout( 10000 )>
```

We also have the ability to set markers in pages. Look at this example:

```
<cfset frapiClass = createObject("java",
"com.intergral.fusionreactor.api.FRAPI")>
<cfset frapi = frapiClass.getInstance()>

<!--- Set a marker --->
<cfset frapi.trace( "Hello World!" )>
```

After this page has run, if we go into FusionReactor and look at the request details for this request we fill see this marker listed along with the time at which it was set:



The ability to set markers within your pages has many useful purposes including following the path of execution, measuring execution time or simply outputting debug information without changing the final output or needing to read through log files.

Example of changing the FusionReactor Configuration via FRAPI

It's possible using FRAPI to change pretty much anything in FusionReactor's configuration. In the example below we toggle an Enterprise Dashboard monitored remote server online/offline programmatically.

In the reactor.conf file the keys "remoting.server.X.*" control the ED monitored servers. The key of interest is "remoting.server.X.online" – you'll have to figure out what "X" is in your case; it's just a serial number. Once you have that, you can use FRAPI to toggle the online state. Here's some CF to do it: *NB you must not have reactor.conf still open in your editor, otherwise the FRAPI code won't be able to change it and you'll get an error, that'll require a restart to clear

If you're in CF and monitoring from within a CFM instance, the following CF code will toggle remoting server 0's state (online/offline)

```
<cfscript>
    frapiClass = createObject( "java",
"com.intergral.fusionreactor.api.FRAPI" );
    frapi = frapiClass.getInstance();
    configuration =
frapi.getRunningConfiguration();
    if( configuration.getStringValue(
"remoting.server.0.online" ) == "true" )
    {
        configuration.setValue(
"remoting.server.0.online", "false" );
    }
    else
    {
       configuration.setValue(
"remoting.server.0.online", "true" );
    }
    frapi.setRunningConfiguration(
configuration );
</cfscript>
```

If you want to change FRAM's configuration you will need to look at FRAM's reactor.conf (/FusionReactor/instance/FRAM/conf/reactor.conf) to get the remoting server number. Then place the following JSP in /FusionReactor/tomcat/webapps/ROOT, and run it as http://localhost:8087/enterprise-toggle.jsp

```
<%@ page
import="com.intergral.fusionreactor.api.*"
%>
<%
    FRAPI frapi = FRAPI.getInstance();
    ConfigurationSurrogate configuration =
frapi.getRunningConfiguration();
    if( configuration.getStringValue(
"remoting.server.0.online" ).equals("true"
) )
    {
        configuration.setValue(
"remoting.server.0.online", "false" );
    }
    else
    {
       configuration.setValue(
"remoting.server.0.online", "true" );
    }
    frapi.setRunningConfiguration(
configuration );
%>
```

Getting the FusionReactor Request ID for the current request

It's possible using FRAPI to get the FusionReactor Request ID of the current running page or request. This can be done by looping over the list of current tracked requests and matching the thread name of the request thread with the thread name on tracked of the FusionReactor request surrogate. The following code example is written in cfscript:

```
<cfscript>
 frapiClass =
createObject("java","com.intergral.fusionr
eactor.api.FRAPI");
 frapi = frapiClass.getInstance();
 requestArray =
frapi.getRunningRequests();
 writeoutput("<h2>Current Request</h2>");
 myThreadName =
createObject("java","java.lang.Thread").cu
rrentThread().getName();
 for (i=1; i LTE ArrayLen(requestArray);
i=i+1)
        {
   if (requestArray[i].getThread() eq
myThreadName)
        {
          writeoutput("ID:
#requestArray[i].getFusionRequestId()#<br/
>");
       }
 }
</cfscript>
```

```
<%@ page
import="com.intergral.fusionreactor.api.*"
%>
<%
    FRAPI frapi = FRAPI.getInstance();
    ConfigurationSurrogate configuration =
frapi.getRunningConfiguration();
    if( configuration.getStringValue(
"remoting.server.0.online" ).equals("true"
) )
    {
        configuration.setValue(
"remoting.server.0.online", "false" );
    }
    else
    {
       configuration.setValue(
"remoting.server.0.online", "true" );
    }
    frapi.setRunningConfiguration(
configuration );
%>
```

Getting the FusionReactor Request ID for the current request

It's possible using FRAPI to get the FusionReactor Request ID of the current running page or request. This can be done by looping over the list of current tracked requests and matching the thread name of the request thread with the thread name on tracked of the FusionReactor request surrogate.

The following code example is written in cfscript:

```
<cfscript>
   frapiClass =
createObject("java","com.intergral.fusionr
eactor.api.FRAPI");
   frapi = frapiClass.getInstance();
   requestArray =
frapi.getRunningRequests();
   writeoutput("<h2>Current
Request</h2>");
   myThreadName =
createObject("java","java.lang.Thread").cu
rrentThread().getName();
   for (i=1; i LTE ArrayLen(requestArray);
i=i+1)
   {
         if (requestArray[i].getThread()
eq myThreadName)
         {
            writeoutput("ID:
#requestArray[i].getFusionRequestId()#<br/
>");
         }
   }
</cfscript>
```

Summary

As we hope you can see, the FRAPI is a very powerful new tool which has a lot of potential. As ever, if you have any feedback or suggestions for this or any other feature within FusionReactor, please let us know and we will try our best to include them in upcoming builds.

FR Administration Manager (FRAM)

FRAM provides a completely isolated instance of FusionReactor which is just for monitoring and instance management, it will continue running even following a JVM crash. FRAM manages and keeps track of all servers

(and instances) so you can quickly and easily install FusionReactor across your environment.

Instance Manager

- Initial Server Scan
- Initial Update
- Add Server & Scan System
- Install FusionReactor Instance
- Uninstall FusionReactor Instance
- Edit FusionReactor Instance
- Edit FusionAnalytics Connector
- Update FusionReactor Instance
- Corrupt Configurations
- Unsupported Configurations

Setting Up FRAM with HTTPS

Instance Manager

The **Instance Manager** (in versions prior to 4.0 only available in the Enterprise Edition) is now available to all editions, and can be found in the **FusionReactor Administration Manager** (FRAM). FRAM is an isolated management instance of FusionReactor, and provides a monitoring instance and management functionality for other instances. For **Enterprise Edition** customers, FRAM also houses an **Enterprise Dashboard** monitor, from where individual instances - including those on other machines - can be monitored.

The Instance Manager makes installing and managing FusionReactor instances as easy as one click of a button. From the Instance Manager you are able to **add**, **remove** and **edit** FusionReactor Instances, as well as setting them up for log shipping to FusionAnalytics.

Some settings are only editable from the Administration Manager, making instance management a central and secure process.



Add Server

In the Instance Manager you are able to search the local system for installed servers. This scan is performed automatically after a FusionReactor 4.0.x installation as covered by the Initial Server Scan Guide. This initial scan only goes as deep as 2 sub directories from the root of the installed filesystem. If this scan did not find all of your installed servers, you can run the scan again and specify your own parameters for the search. For more information please read the Add Server & Scan System documentation.

The Add Server & Scan System page also covers how to manually add an installed server to the Instance Manager by providing the server root directory.

Manage Instances

This section of the Instance Manager list all of the servers and FusionReactor instances known to the Administration Manager.

The Instance Manager displays information as well as provide management actions for each FusionReactor instance listed in the table of server instances that the Instance Manager knows of.

Instance Information



Figure 1: Table Headers

Figure 1: The information available for each FusionReactor instance.

Field Name	Description
Instance	This column will display the actions available for each instance, see below section **Available Actions** for more information on the actions available for the server instances.
Instance Link	If FusionReactor has been installed into this instance, this column contains a link to that instance. If not, this column contains a "-" character.
Port	If FusionReactor has been installed into this instance, this column contains the port of the installed FusionReactor instance. If not this column will contain a "-" character.
Version	If FusionReactor has been installed into this instance, this column contains the version of FusionReactor currently installed into the server instance. If not this column will contain a "-" character.
	This column contains actions to be performed on multiple instances. The action will be performed on every instance that is selected (by checking the box in this column for each instance), or by clicking the hyperlink *"**Select All**"* to select all of the instances listed in the Instance Manager. For a list of available actions see the below section **Available Actions.**

Available Actions

For each server listed in the Instance Manager the following action is available:

Button	Name	Description

	Remove Server	If no instances within a server have FusionReactor installed you can remove it from the Instance Manager. **NOTE** This option just removes it from FusionReactor. It does **not** rem ove the server from the machine. The server can be added again later using the **scan** form at the top of the page.
	Remove Server Unavailable	If there are at least one server instance with FusionReactor installed you cannot remove the server from the Instance Manager. To do this firstly remove all of the FusionReactor instances installed on instances within this server. Please read the Uninstall FusionReactor Instance Guide on how to remove a FusionReactor instance.

For each server instance listed in the Instance Manager the following action is available:

Button	Name	Description
	Install FusionReactor Instance	Clicking this button will install FusionReactor into the selected server instance, as well as creating a FusionReactor instance listed in the Instance Manager for easy access. See the Install FusionReactor Instance document ation for more information on how to install a FusionReactor instance.

	Install FusionReactor Instance Unavailable	This button will be disabled if there is already a FusionReactor instance installed into the selected server instance. See the below section for actions available for each FusionReactor instance.
	Select this Instance	Tick this check box if you wish to perform any of the actions located on the table header (See the above section **Instance Information**) on all of the selected instances. All of the server instances listed can be selected by clicking the hyperlink *"**Select All**"*.

For each FusionReactor instance listed in the Instance Manager the following actions are available:

Button	Name	Description
	Uninstall FusionReactor Instance	If FusionReactor has been installed into a server instance, clicking this button will uninstall FusionReactor from the selected instance. The Uninstall FusionReactor Instance documentation covers this in more detail.
	Edit FusionReactor Instance	Clicking this button will enable you to edit the web server properties of the selected FusionReactor instance. See the online documentation on Edit FusionReactor Instance for more detailed information.
	Add to Dashboard	The selected FusionReactor instance can be easily monitored by the **Enterprise Dashboard** available in the FusionReactor Administration Manager (FRAM) if you have installed an Enterprise License. Clicking this button will add the selected instance to FRAM's Enterprise Dashboard. **Please note that the Enterprise Dashboard is only available in the Enterprise version of FusionReactor.** For more information see the Enterprise documentation.

	Edit FusionAnalytics Connector	Clicking this button will allow you to set up the selected instance to transfer logs to FusionAnalytics, as well as modify the logging settings for this instance. The Edit FusionAnalytics Connector documentation covers setting up an instance for log transfer in more detail.
	Update FusionReactor Instance	If the server instance has an older version of FusionReactor installed than that installed on the FusionReactor Administration Manager (FRAM), the instance can be easily updated to run the same FusionReactor version as FRAM directly from the Instance Manager. See the online documentation Update FusionReactor Instance on how to update an outdated FusionReactor Instance.

Initial Server Scan

After using the FusionReactor 4.0.x setup to perform a fresh install of FusionReactor, you will (by default) be taken to the FusionReactor Administration Manager (FRAM) and the Initial Server Scan wizard will pop up.

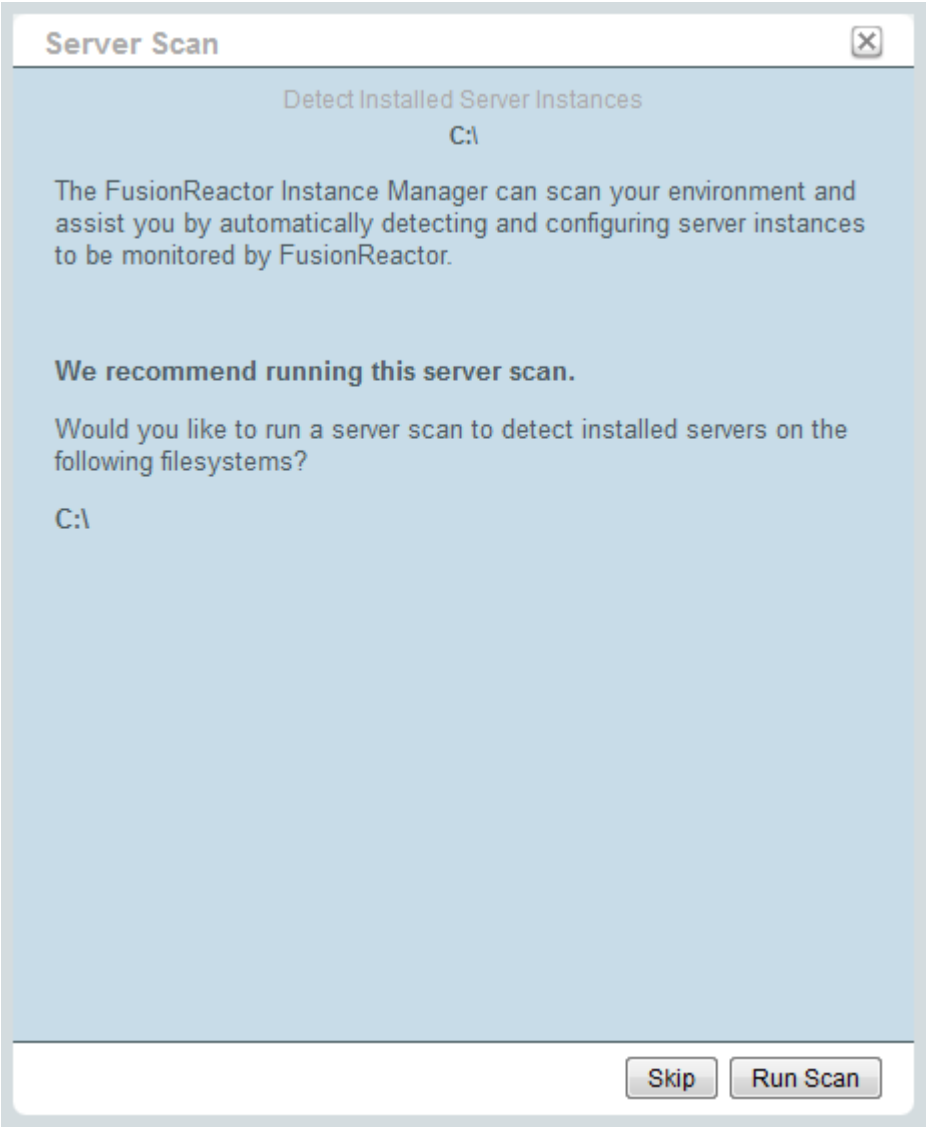


Figure 1: Server Scan Information

Figure 1: Clicking *'Run Scan'* will search your system for servers to monitor. **The initial server scan will only search 2 sub directories deep from the displayed filesystem root.**

If you click 'Skip', the scan will not run and no servers will be discovered by the FusionReactor Administration Manager. Clicking *'Skip'* will take you back to the Instance Manager. You can manually run the server scan with provided parameters from the Instance Manager. Follow the online guide Add Server & Scan System on how to do this.

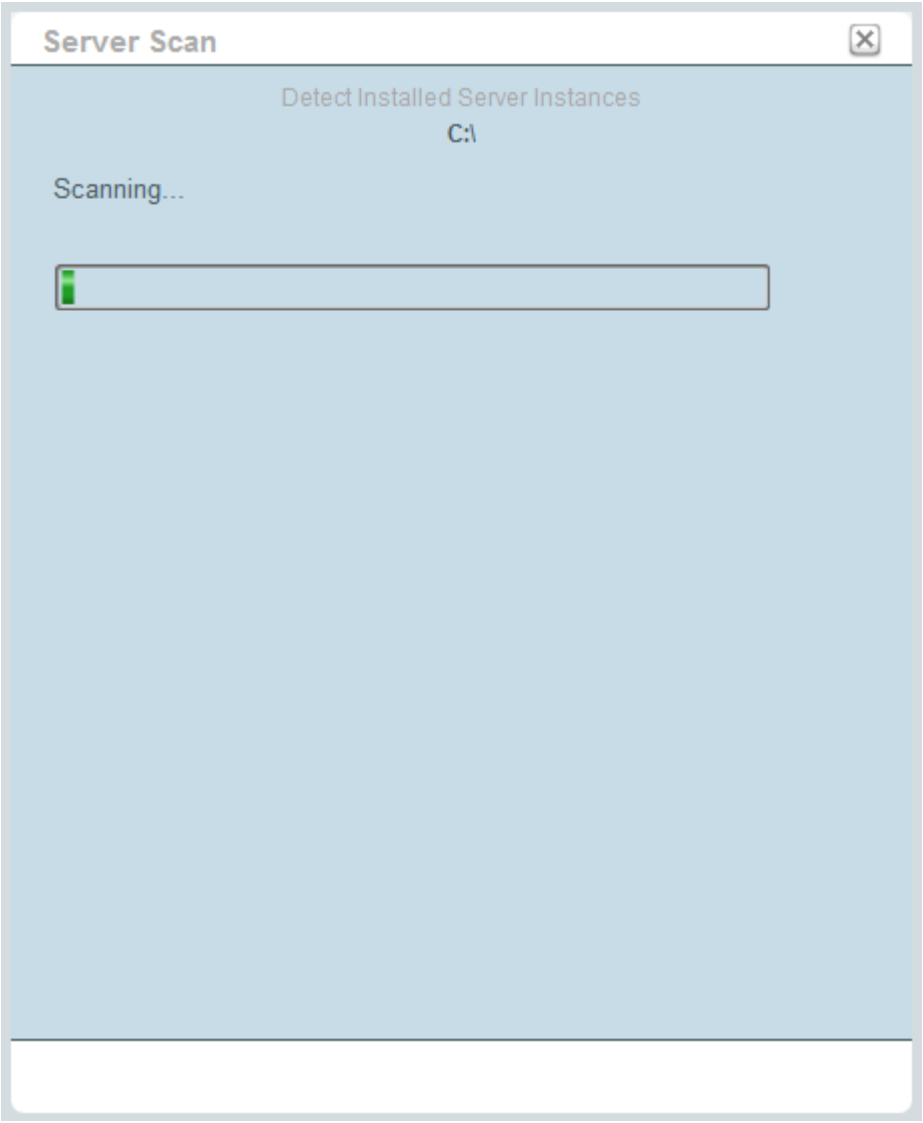


Figure 2: Scanning

Figure 2: The scan is searching the local filesystem for servers that can be monitored by FusionReactor. This scan may take several minutes to complete depending on the size of your hard drive and the number of files stored on it.

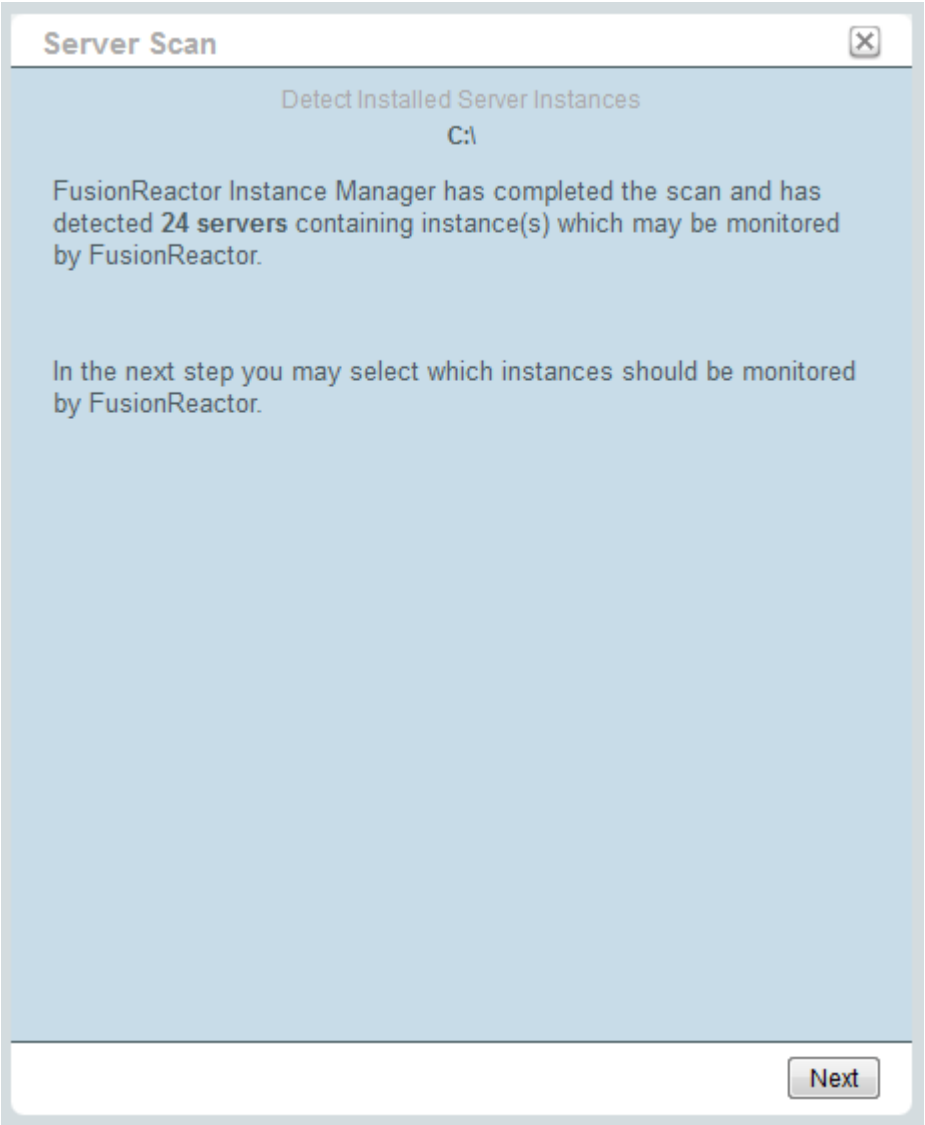


Figure 3: Server Scan Finished

Figure 3: Once the scan has completed, all detected servers that can be monitored by FusionReactor will be displayed.

Click *'Next'* to see a complete list of the servers which were discovered.

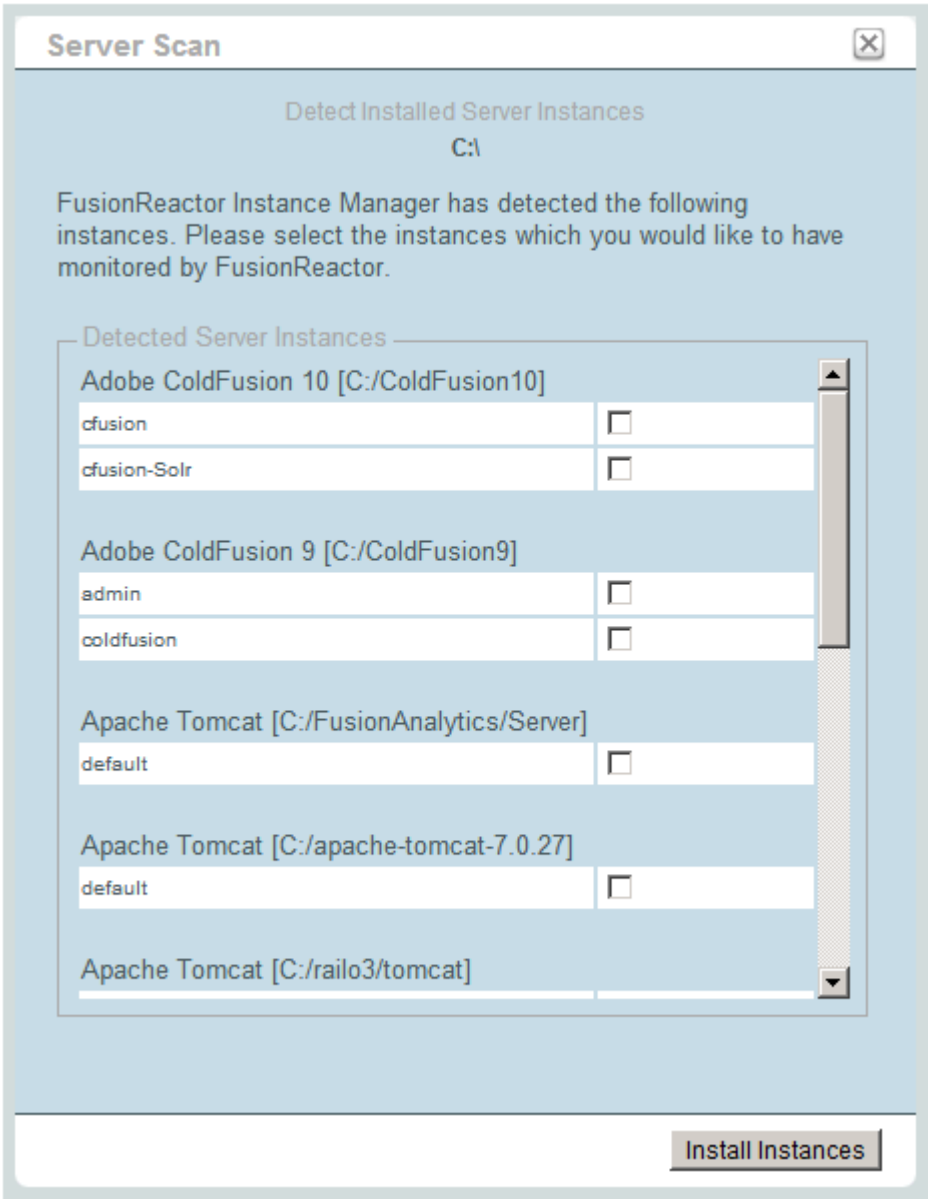


Figure 4: Select Server Instances

Figure 4: Select the server instances into which you wish to install FusionReactor by ticking the check box associated with each server instance.

Clicking *'Install Instances'* with one or more servers instances checked will install FusionReactor into each one of the selected instances. More detailed information is available on the subject of installing FusionReactor Instances.

If you do not wish to install FusionReactor into any server instances at this point then either click the "X" button at the top right of the dialog box, or click *'Install Instances'* with no check boxes selected. This will take you back to the Instance Manager where you also have the opportunity to add FusionReactor instances by following the online guide Install FusionReactor Instance.

Initial Update

After using the FusionReactor 4.0.x setup to perform an update from FusionReactor 3.x, you will (by default) be taken to the FusionReactor Administration Manager (FRAM) and the Initial Update wizard will pop up.

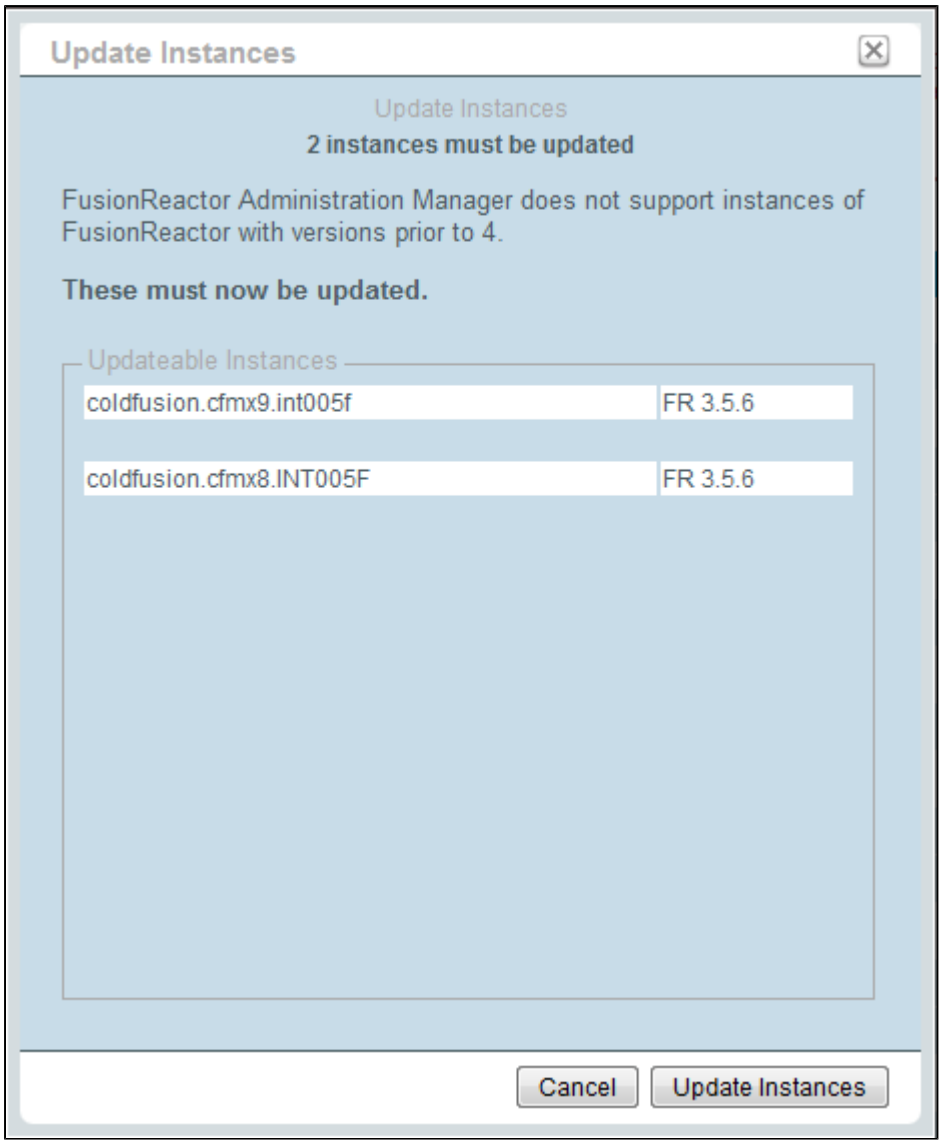


Figure 1: Instances that can be updated

Figure 1: You are shown a list of all the instances that can be updated to the latest version of FusionReactor. This list contains the FusionReactor instance name and the version of FusionReactor currently installed into these instances.

Click *'Update Instances'* to continue and update all of the listed instances. The Instance update process is covered by the online documentation Update FusionReactor Instance.

Click *'Cancel'* if you do not wish to update all of the listed instance at this point.

> It's important to note that your original 3.x instances might continue to run, but they should be updated as soon as possible. FusionReactor doesn't support a mixed configuration of 4.x and pre-4.x instances.

Add Server & Scan System

Adding a Server

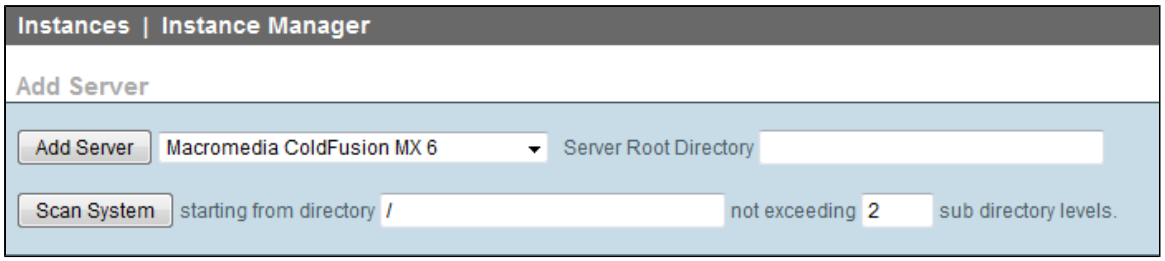


Figure 1: Add Servers

To add a new server to the Instance Manager, select the type of server you want to add from the drop-down box and then enter the file path to the root directory of the server. You must enter the full path, for example `C:\CFusionMX`

Clicking the *'Add Server'* button will cause FusionReactor to validate the server. If the file directory entered contains a valid server the specified server will appear in the instance list under the **Manage Instances** section (See Figure 2) along with any other identified servers and instances.

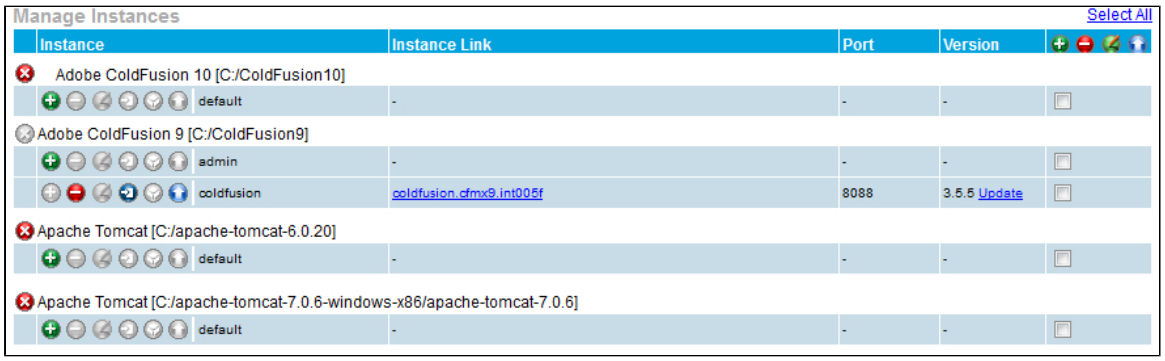


Figure 2: Manage Instances

Scanning the System for New Servers

> **Improved server detection**
>
> Server detection for both Coldfusion and Railo Servers has been improved in 4.5.0

FRAM can scan the system for new servers and server instances. You can specify a custom starting point by entering a full directory path into the *'starting from directory'* field, or you can leave it with its default value which will search from the root of the filesystem the FusionReactor Administration Manager (FRAM) was installed into. (It will not scan multiple drives, for example, on Windows, it will not scan `C:` then `D:` and then `E:`)

You can limit the search by specifying how many folders deep FusionReactor should search through. (See Figure 1)

> **NOTE**
>
> Please be aware that asking FusionReactor to search more than 3 or 4 levels deep on a complex file system can take several minutes to complete.

If FusionReactor Administration Manager can't detect your instance, you will have to install it manually (See Manual Instance Installation).

Install FusionReactor Instance

Inside the Instance Manager, you will be able to install FusionReactor into any server instance listed. To do this,

click on the button next to the desired server or click the button in the rightmost column header to install FusionReactor into one or more server instances (you need to tick the check boxes of the desired server instance or click the hyperlink *"Select All"* if you use the button in the column header).

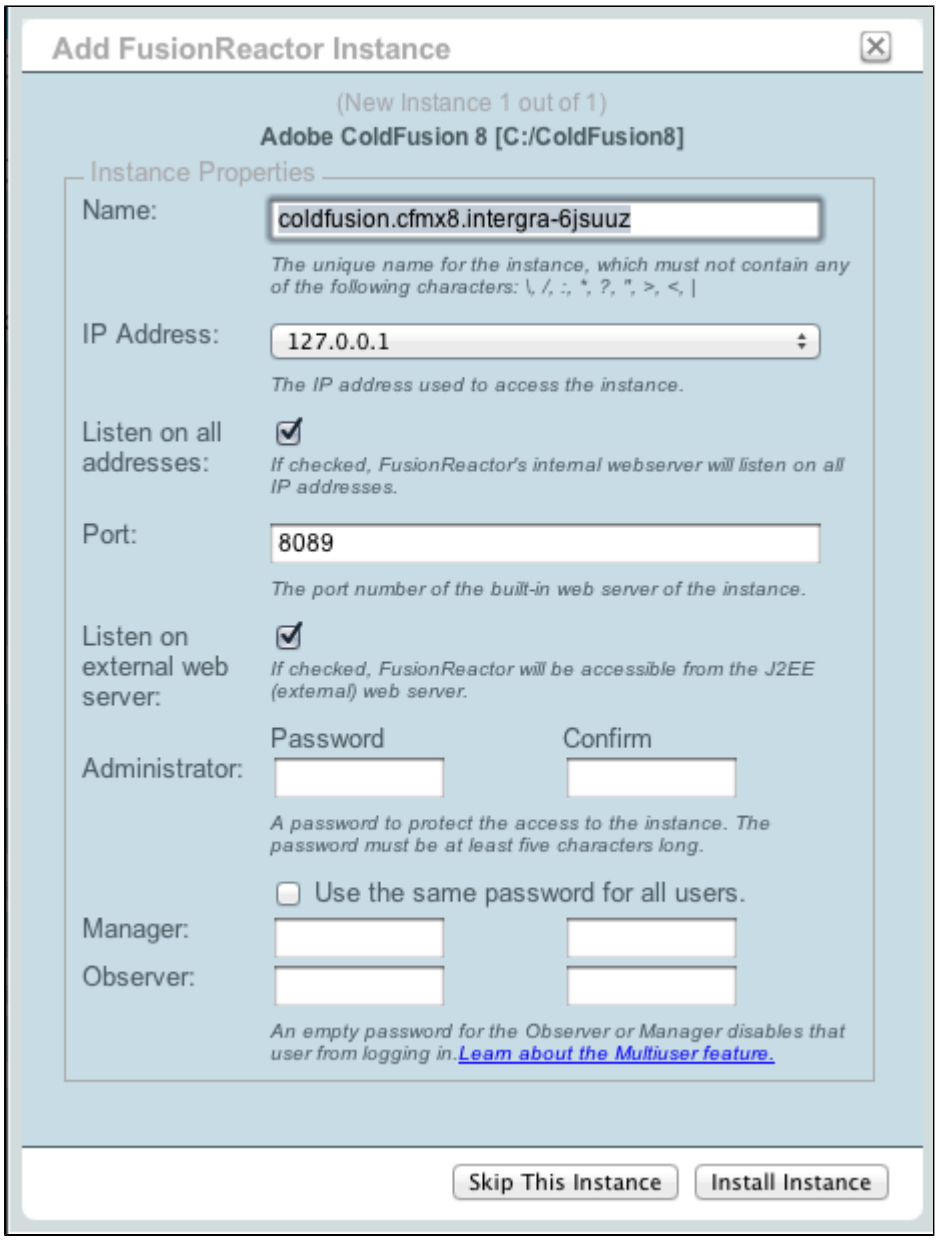


Figure 1: Configure the FusionReactor Instance

Figure 1: For each server instance you will need to provide configuration information used to set up FusionReactor.

Configuration	Description
Name	This is the unique name of the FusionReactor instance to be installed. The name must **not** contain any non-latin characters, or any of the following: \, /, :, *, ?, ", >, <, |

IP Address	This is the IP address used to access the FusionReactor instance on the internal web server provided by FusionReactor. The addresses listed are the ones available on the machine local to the FusionReactor Administration Manager (FRAM). N.b. FusionReactor is also accessible using your external web server, on the url * http://localhost/fusionreactor* ;
Listen on all addresses	Tick this box if the internal web server provided by FusionReactor should listen on all IP addresses on all interfaces, not just that named in the previous input box.
Port	This is the port number used to access the FusionReactor instance using FusionReactor's internal web server.
Listen on external web server	Tick this box if you want FusionReactor to also listen on your J2EE server's web server.
Administrator Password	Provide a password for the **Administrator** user on the new instance. For more information about the user roles in FusionReactor see Types Of User Account
Use the same password for all users	Tick this check box if you wish to use the Administrator password for all user roles. Selecting this option will disable the Manager and Observer Password fields. For more information about the user roles in FusionReactor see Types Of User Account
Manager Password	Provide a password for the **Manager** user on the new instance. For more information about the user roles in FusionReactor see Types Of User Account
Observer Password	Provide a password for the **Observer** user on the new instance. For more information about the user roles in FusionReactor see Types Of User Account

The instance initially listens on the internal webserver (this can be changed in XXX), and - optionally - the external web server.

Once you have correctly provided all of the configuration details click *'Install Instance'* to install FusionReactor into the selected server instance and create a FusionReactor instance with the specified details.

Click *'Skip This Instance'* to cancel the installation of this FusionReactor instance. If you have selected to install more than one instance, clicking *'Skip This Instance'* will advance the wizard to the next instance.

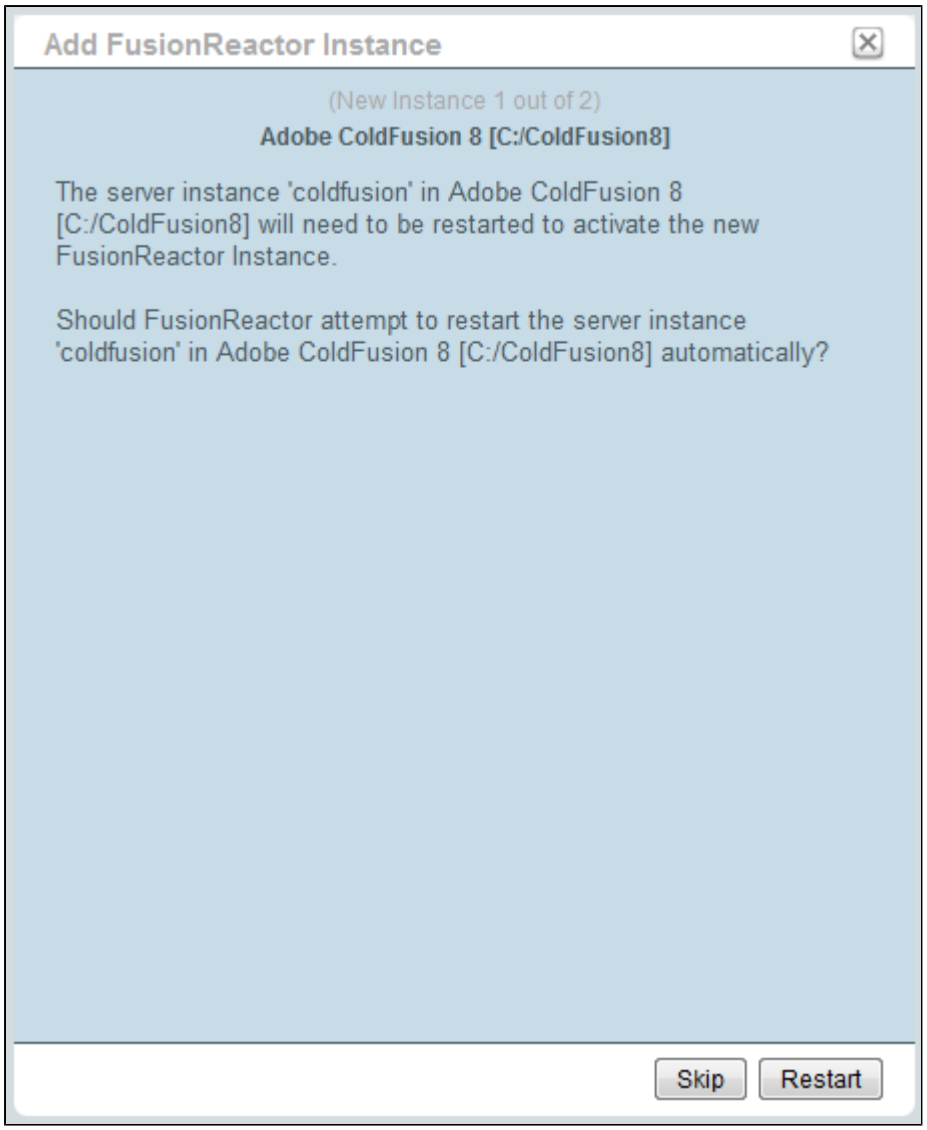


Figure 2: Restart the Server

Figure 2: To start FusionReactor on the server instance, the server instance must be restarted. By clicking *'Rest art'* The Instance Manager can attempt to restart the server instance automatically (See Figure 3).

If you wish to restart the server at another point, or wish to manually restart it, click *'Skip'* (See Figure 5)

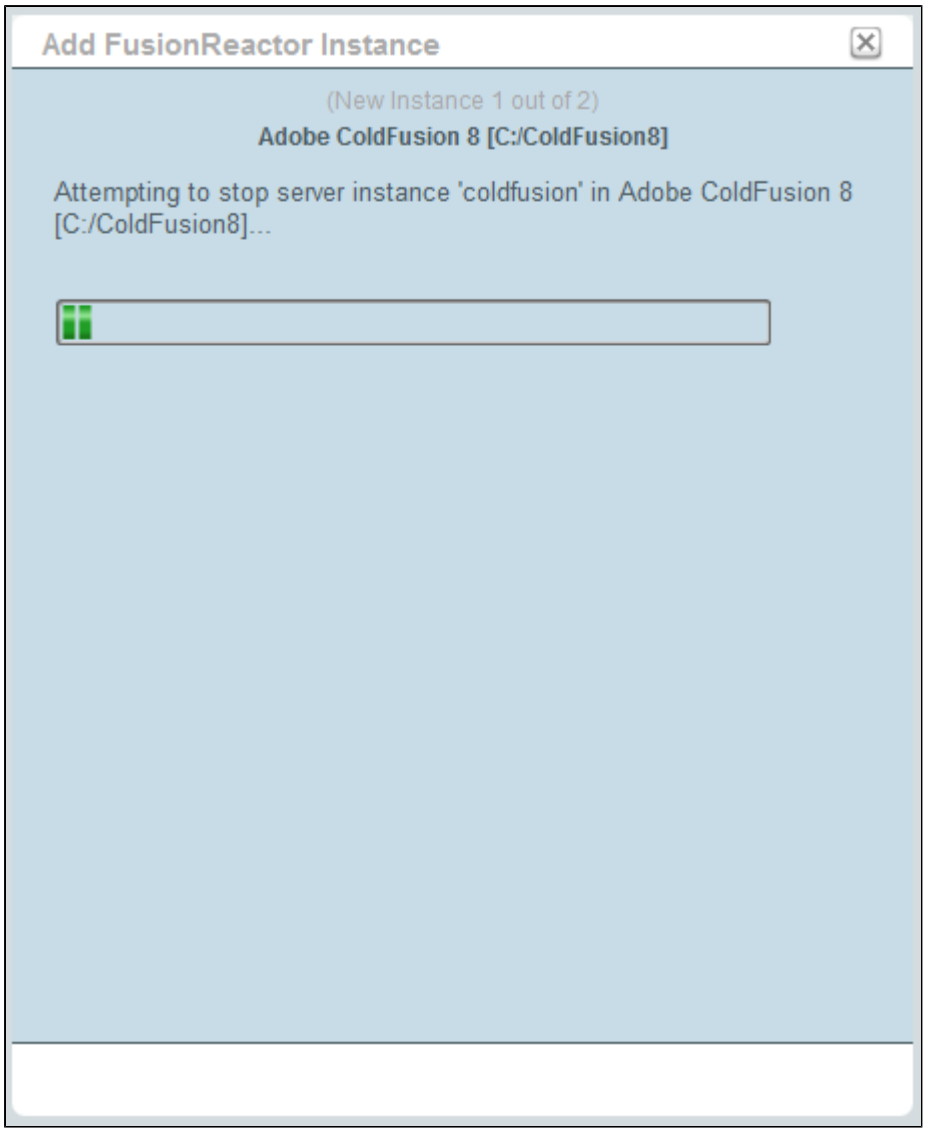


Figure 3: Restart in Progress

Figure 3: FusionReactor is attempting to stop the server instance. Please note that stopping the server may take several minutes depending on your system configuration.

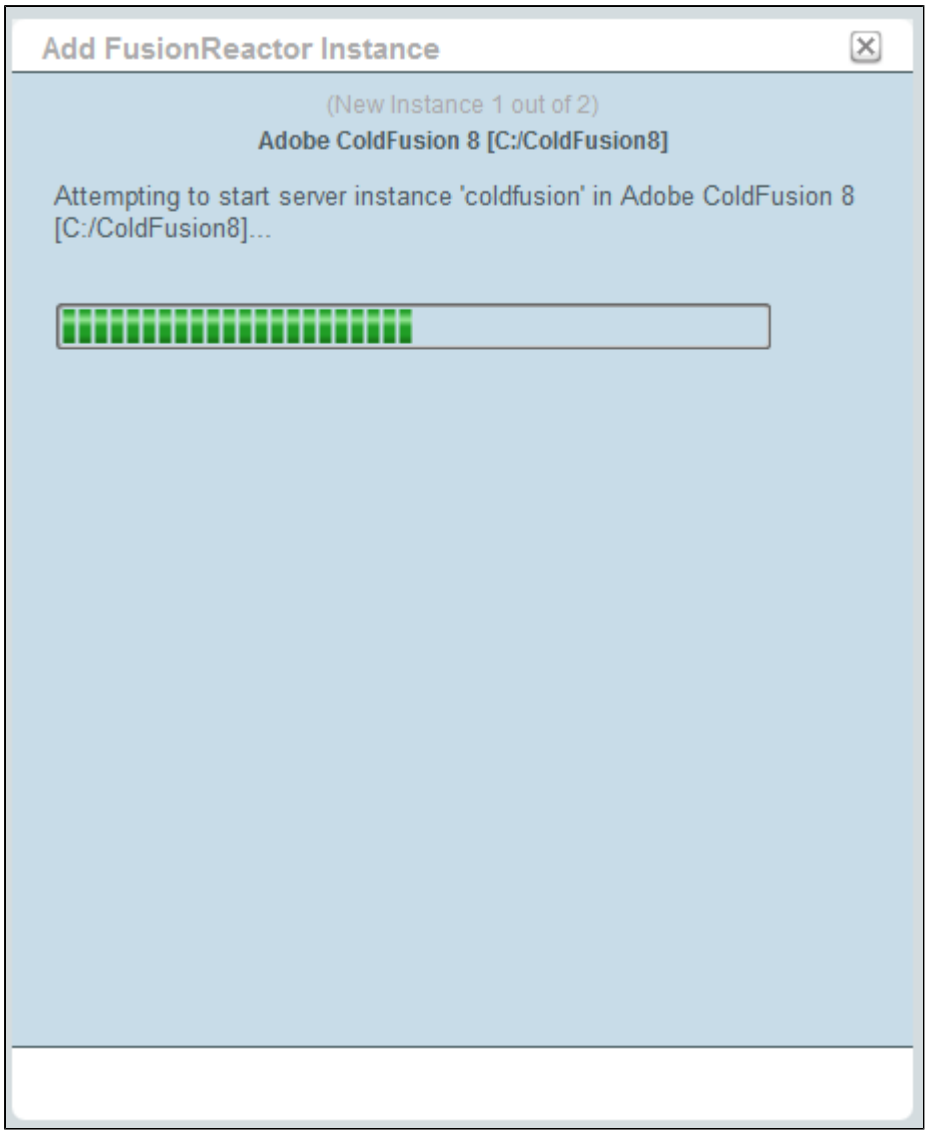


Figure 4: Restart in Progress

Figure 4: FusionReactor is attempting to start the server instance. Please note that starting the server may take several minutes depending on your system configuration.

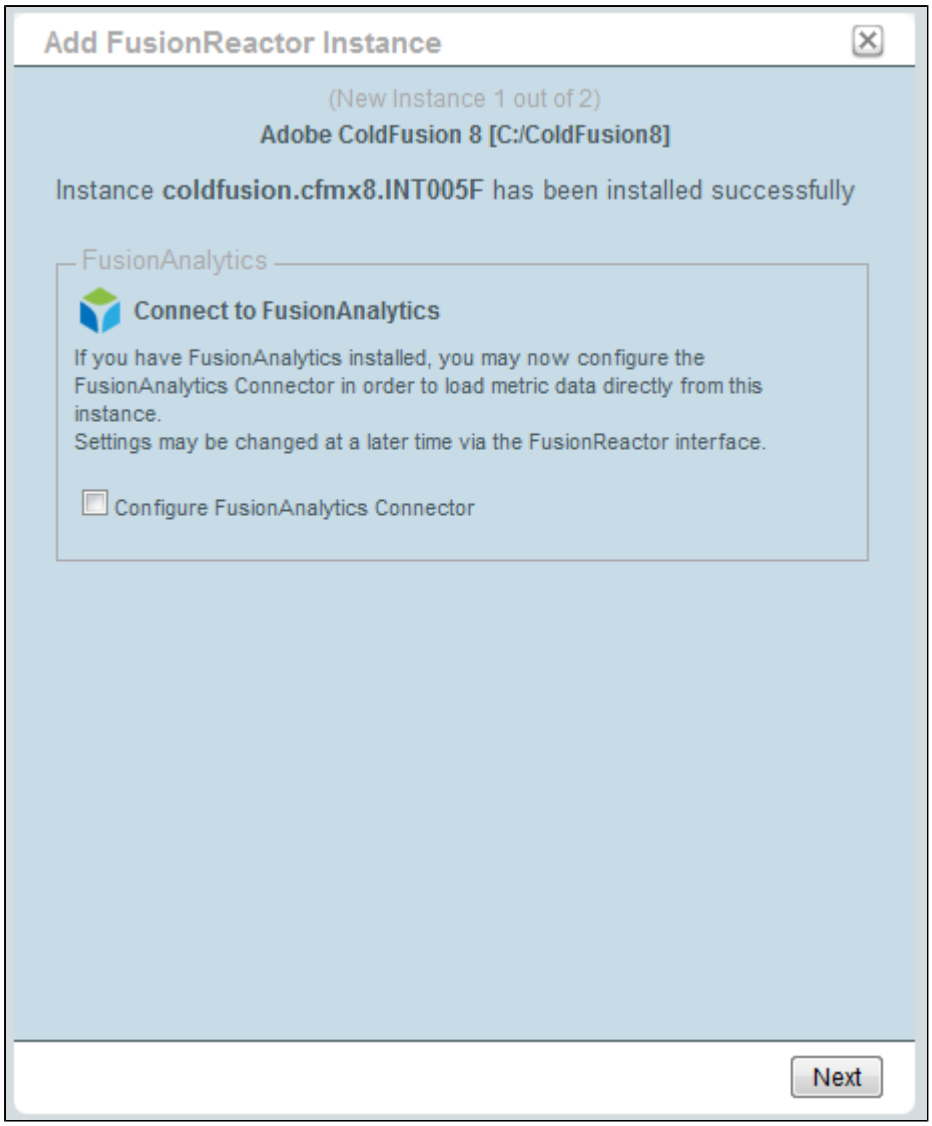


Figure 5: Instance Installation Successful

Figure 5: Once the installation is complete you will have the choice to configure the FusionAnalytics Connector. Tick the check box *'Configure FusionAnalytics Connector'* and click *'Next'* if you wish to configure log transfer to FusionAnalytics now. Otherwise leave the check box unticked.

If you have selected to install multiple FusionReactor instances clicking *'Next'* will advance the wizard to install the next instance.(Go back to Figure 1).

If only one instance was selected you have finished the Installation wizard as displayed by Figure 6.

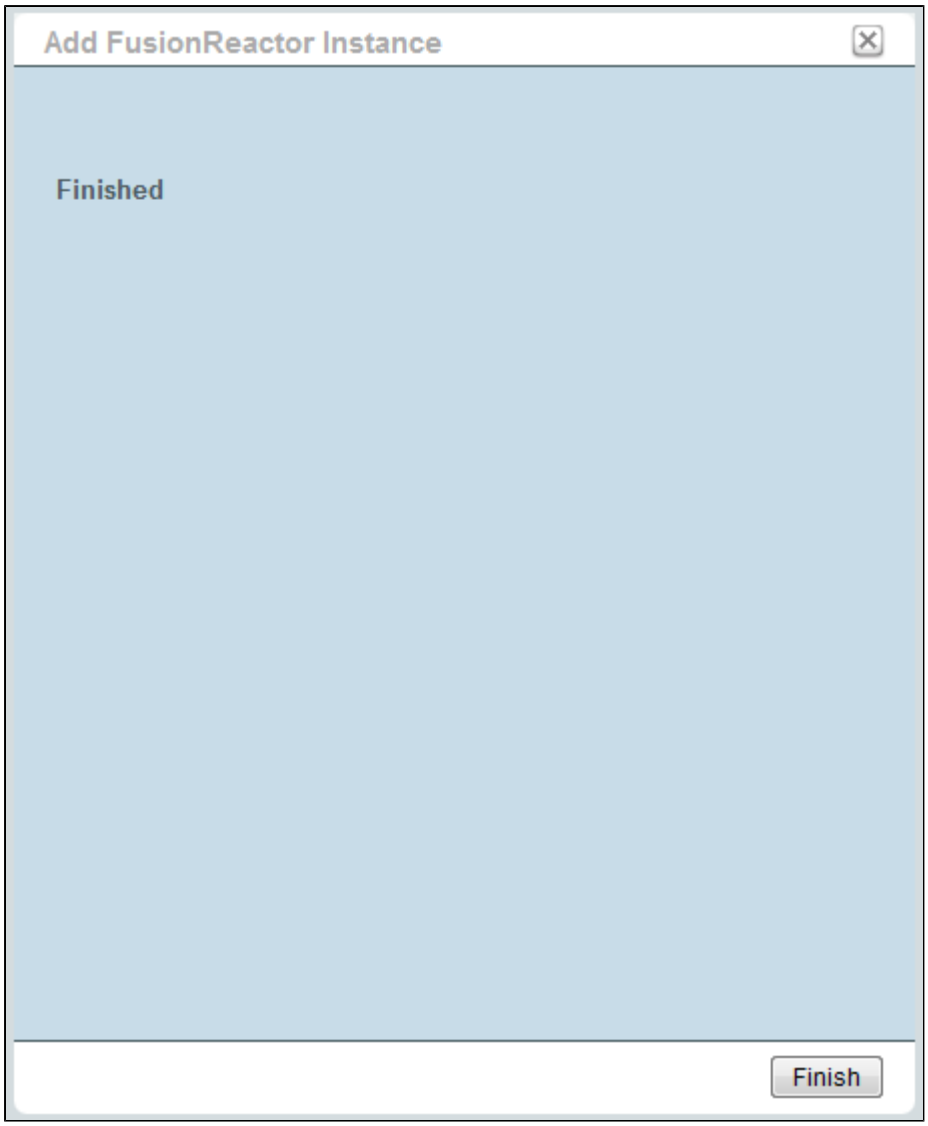


Figure 6: Instance installation finished

Figure 6: FusionReactor has been installed into the selected server instances. Click *'Finish'* to return to the Instance Manager.

Uninstall FusionReactor Instance

Inside the Instance Manager, you will be able to uninstall FusionReactor from any server instance listed. To do this, click on the button next to the desired FusionReactor instance or click the button in the far right column header (you need to tick the check boxes of the desired server instances or click the hyperlink *"Select All"* if you use the button in the column header).

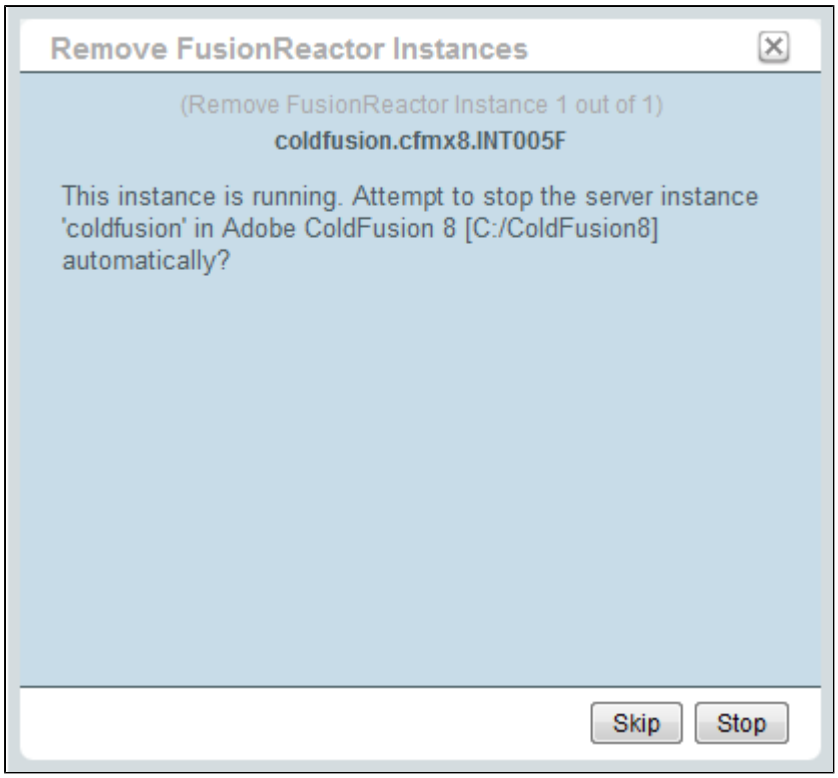


Figure 1: Stop the Server Instance

Figure 1: If the server instance you are attempting to remove FusionReactor from is currently running, it needs to be stopped in order to completely remove the FusionReactor instance.

By clicking *'Stop'* the Instance Manager can attempt to stop the server instance automatically (See Figure 2).

If you wish to stop the server at another point, or wish to manually restart it, click *'Skip'*. If you have selected to uninstall more than one FusionReactor instance, the wizard will advance to the next FusionReactor instance. Otherwise the Finish screen (Figure 6) will be displayed.

> To completely remove the FusionReactor Instance and all of its components the server instance must be completely shut down.

If the server instance you are removing FusionReactor from is not running Figure 3 will be displayed.

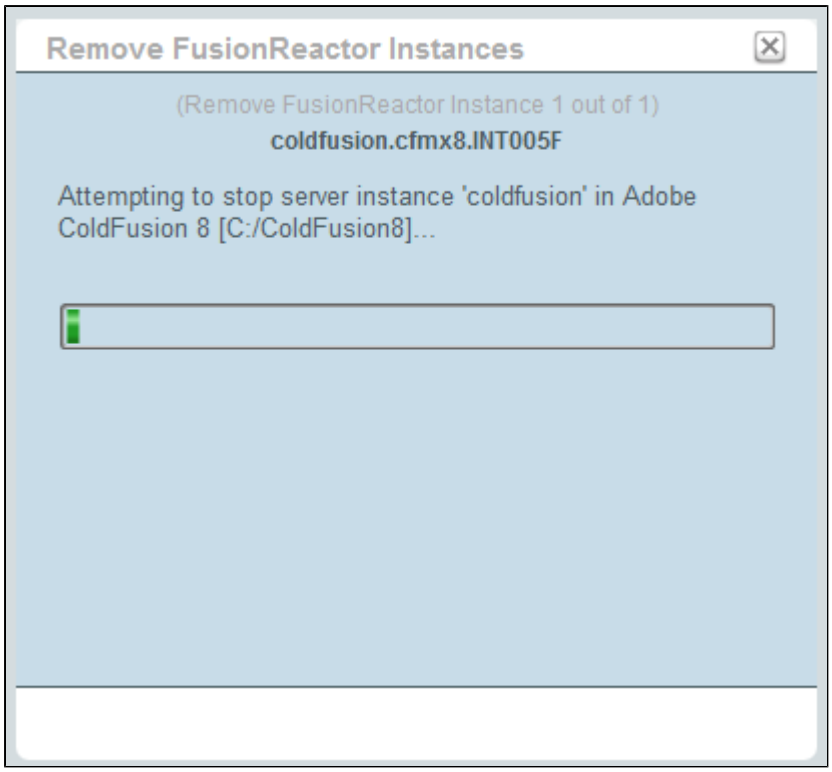


Figure 2: Stop in Progress

Figure 2: FusionReactor is attempting to stop the server instance. Please note that stopping the server may take several minutes depending on your system and server.

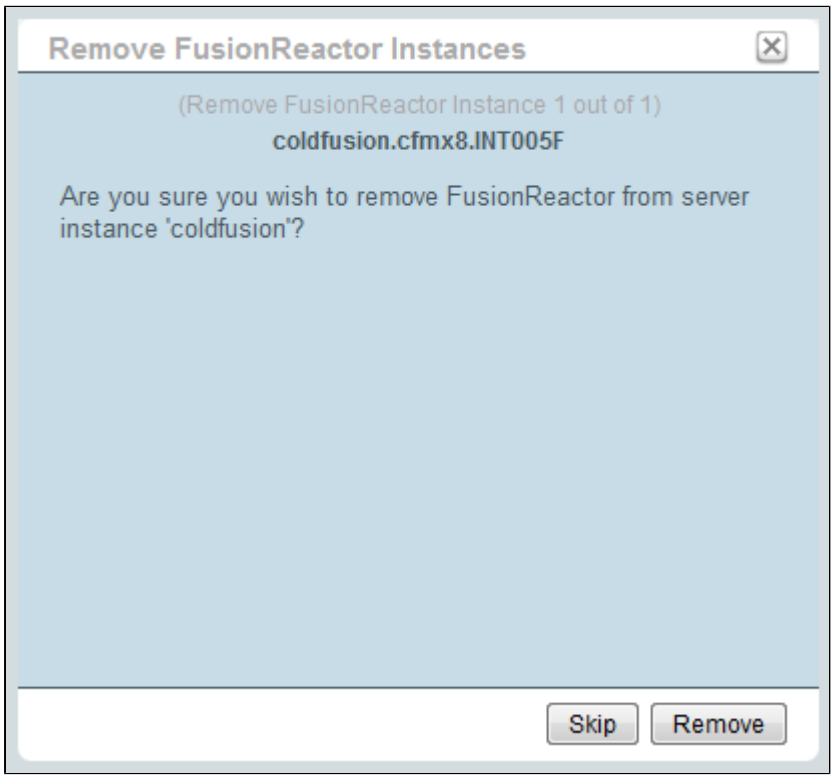


Figure 3: Confirm Uninstall

Figure 3: By clicking *'Remove'* FusionReactor will be completely removed from the server instance.

If you do not wish to uninstall FusionReactor from this server instance press *'Skip'*. If you have selected to remove more than one FusionReactor Instance the wizard will advance to next FusionReactor instance, if not

the finish screen will be displayed (See Figure 7).

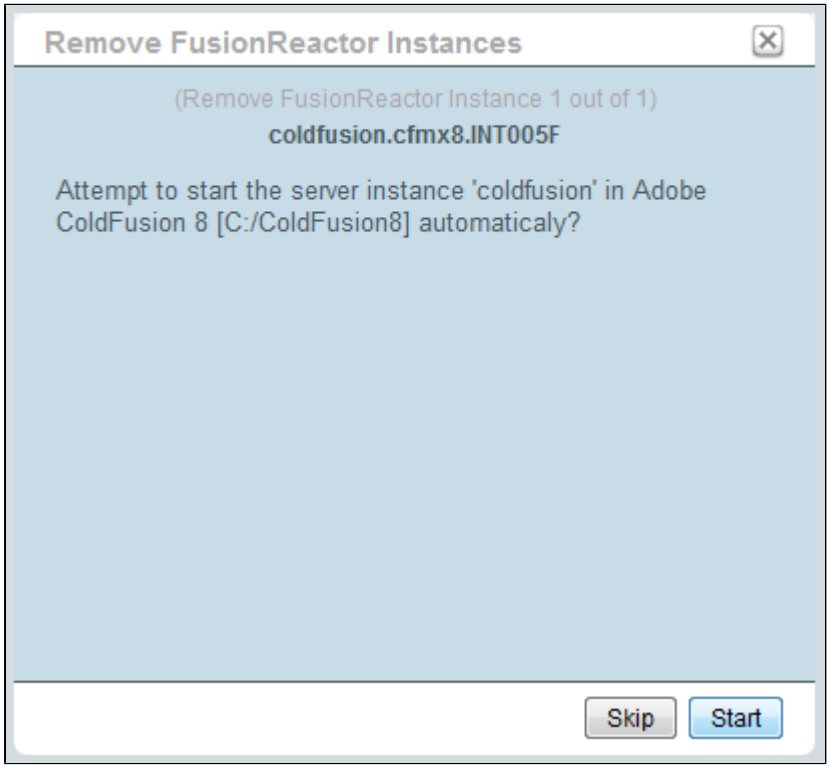


Figure 4: Restart in Progress

Figure 4: By Clicking *'Start'* the Instance Manager will attempt to automatically start the server instance again (See Figure 5).

If you wish to start the server at another point, or wish to manually start it, click *'Skip'* (See Figure 6).

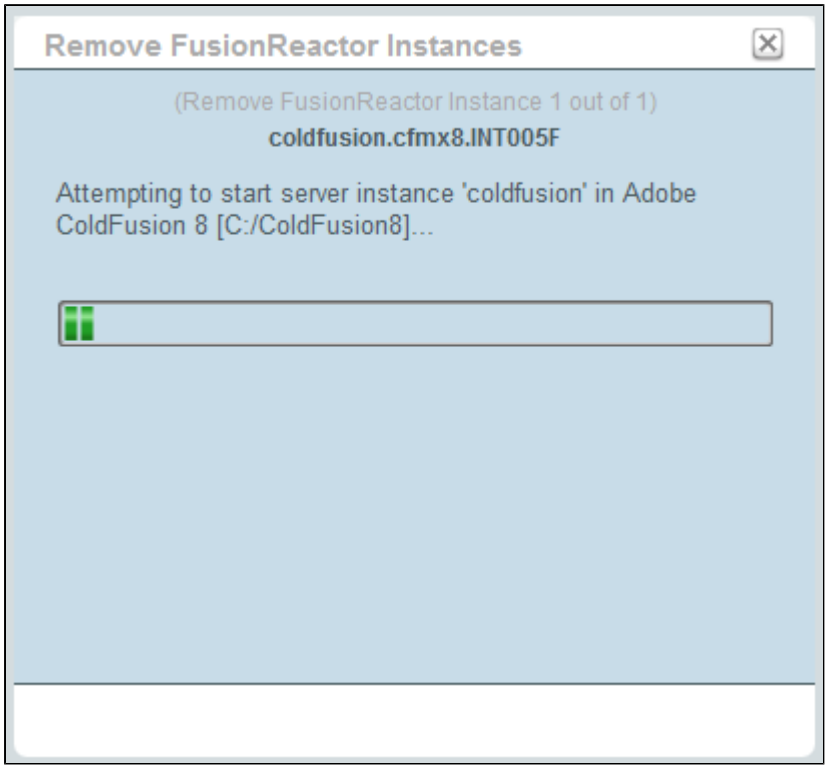


Figure 5: Restart in Progress

Figure 5: FusionReactor is attempting to start the server instance. Please note that starting the server may take several minutes depending on your system configuration.

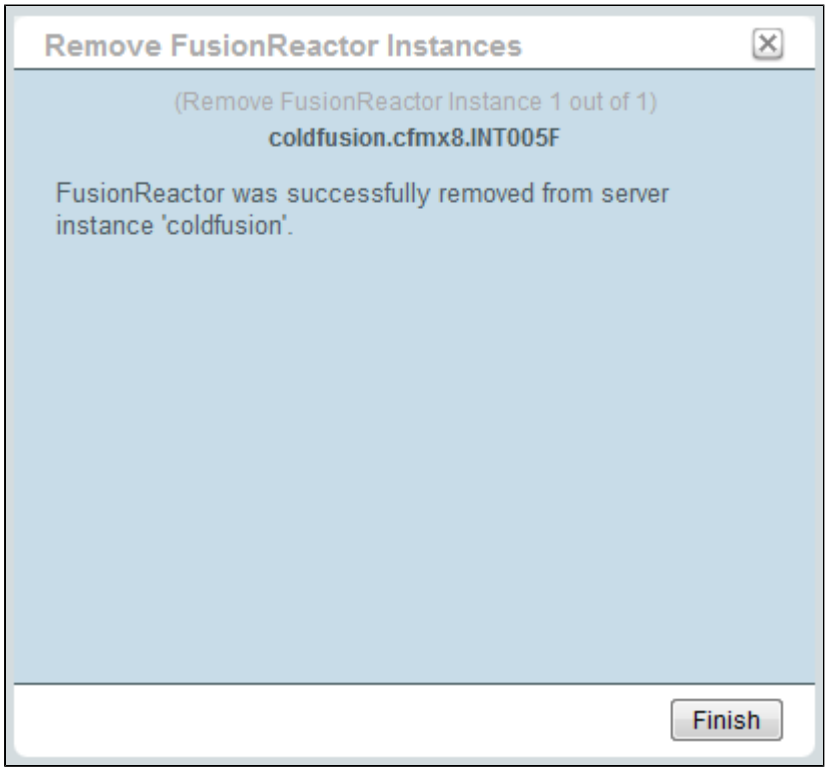


Figure 6: FusionReactor Instance Removed

Figure 6: Once the FusionReactor instance has been removed you will see a confirmation screen where you will be told whether the removal was successful.

If you have selected to only uninstall one FusionReactor instance click *'Finish'* to return to the Instance Manager.

If you have selected to remove more than one FusionReactor instance click *'Next'* to advance the wizard to the next FusionReactor instance (See Figure 1).

If this is the last instance to be removed clicking *'Next'* will display the finish screen as displayed in Figure 7.

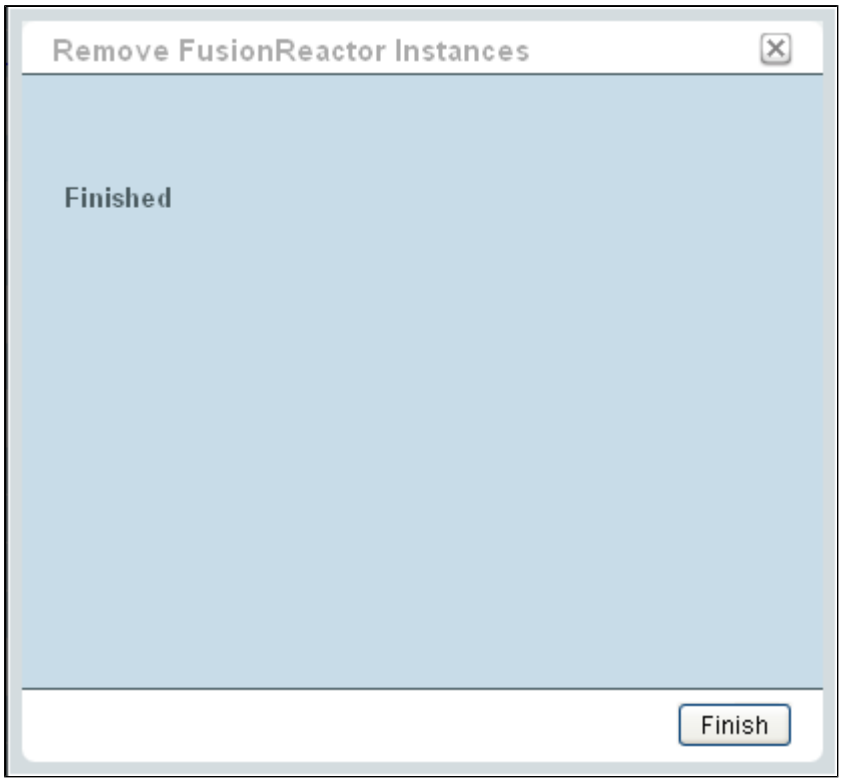


Figure 7: FusionReactor Instance Removed

Figure 7: FusionReactor has been uninstalled from the selected server instances. Click *'Finish'* to return to the Instance Manager.

Edit FusionReactor Instance

Inside the Instance Manager, you will be able to edit any FusionReactor Instance listed. To do this, click on the button next to the desired FusionReactor instance or click the button in the rightmost column header (You need to tick the check boxes of the desired FusionReactor Instances or click the hyperlink *"Select All"* if you use the button in the column header).

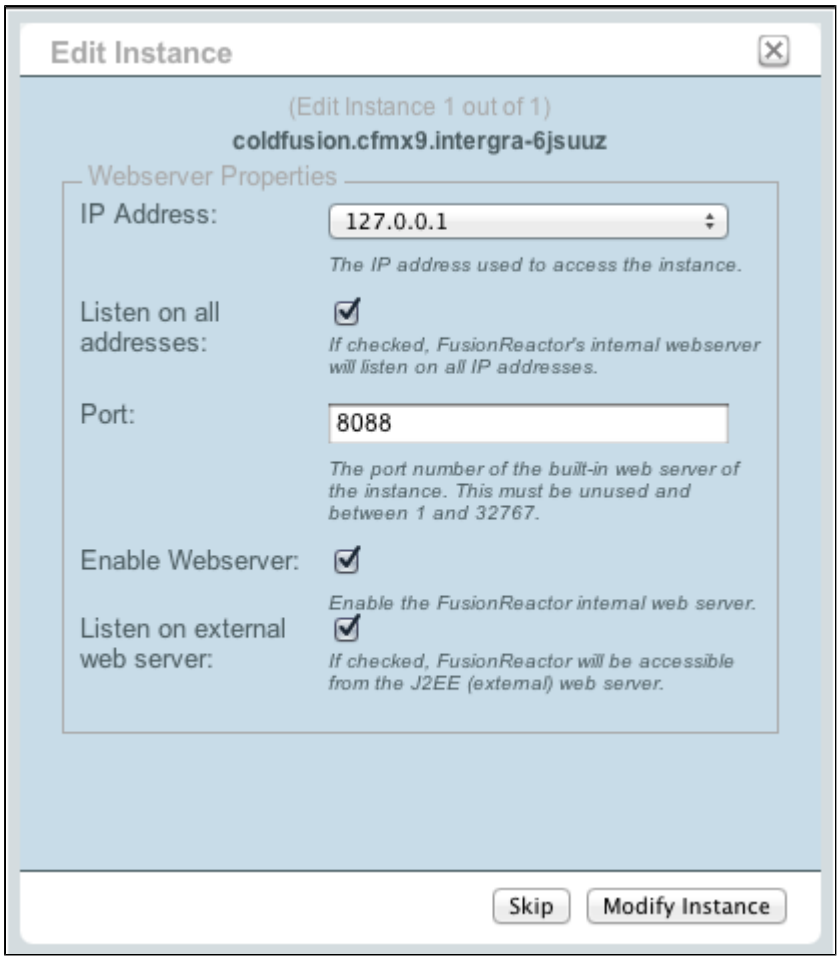


Figure 1: Edit Instance Properties

Figure 1: For each FusionReactor Instance you can edit the web server configuration as displayed in Figure 1.

Configuration	Description
IP Address	This is the IP address used to access the FusionReactor instance on the internal web server provided by FusionReactor. The addresses listed are the ones available on the machine local to the FusionReactor Administration Manager (FRAM).
Listen on all addresses	Tick this box if the internal web server provided by FusionReactor should listen on all of the IP addresses listed in the above field. This makes the FusionReactor instance available on all of the listed addresses.
Port	This is the port number used to access the FusionReactor instance using FusionReactor's internal web server.
Enable Webserver	Tick this check box if you want to enable FusionReactor's internal webserver using the above properties.

Listen on external web server	Tick this box if you want FusionReactor to also listen on your J2EE server's web server.

N.b. At least one web server channel must be enabled by ticking "Enable Webserver" or "Listen on external web server" (or both). Failing to ensure at least one box is ticked will produce a warning dialogue, and the settings won't be submitted.

Once you have correctly provided all of the configuration details click *'Modify Instance'* to edit the FusionReactor Instance.

Click *'Skip'* to cancel the edit of this FusionReactor instance. If you have selected to edit more than one instance, clicking *'Skip'* will advance the wizard to the next instance. Otherwise the finish screen (Figure 3) will be displayed.

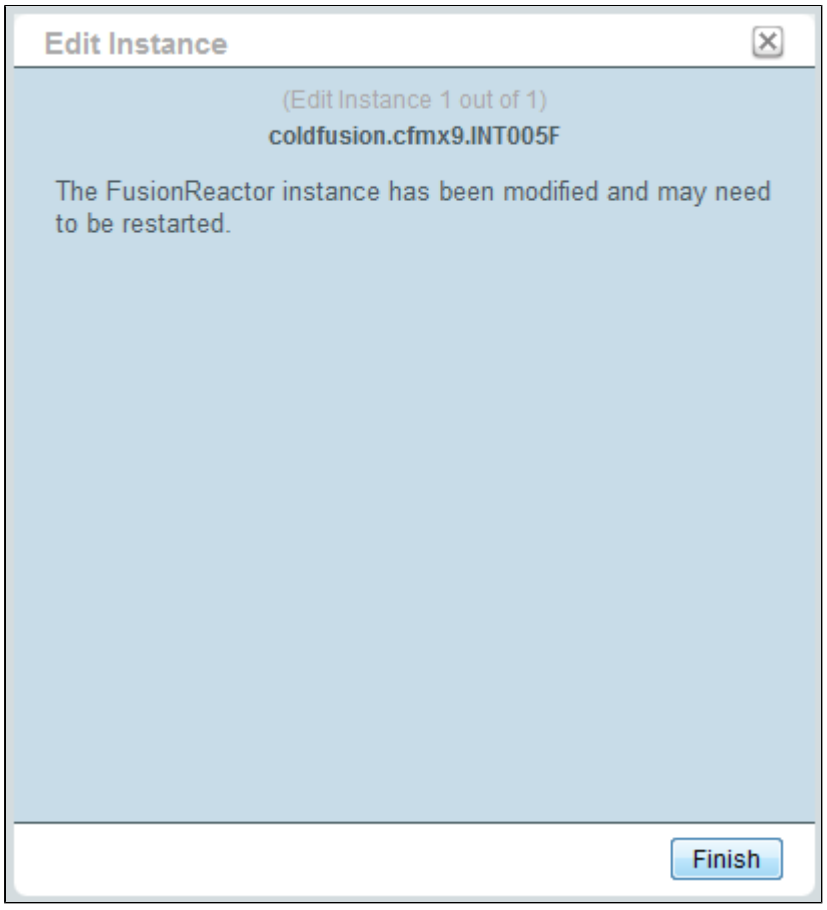


Figure 2: Restart Warning

Figure 2: This screen is just a reminder to let you know that certain configuration changes, such as new web server details, will not be applied until the instance is restarted.

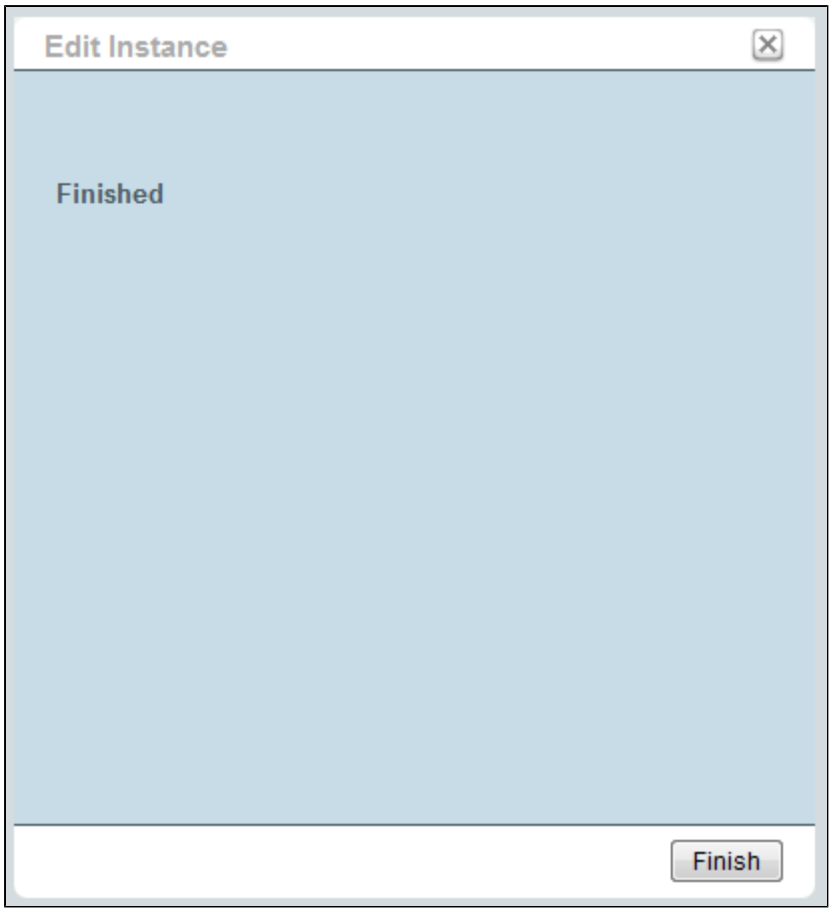


Figure 3: Finished

Figure 3: The FusionReactor Instance has been modified. Click *'Finish'* to return to the Instance Manager.

Edit FusionAnalytics Connector

Inside the Instance Manager, you will be able to edit the FusionAnalytics Connector for any FusionReactor Instance listed. To do this, click on the button next to the desired FusionReactor Instance.

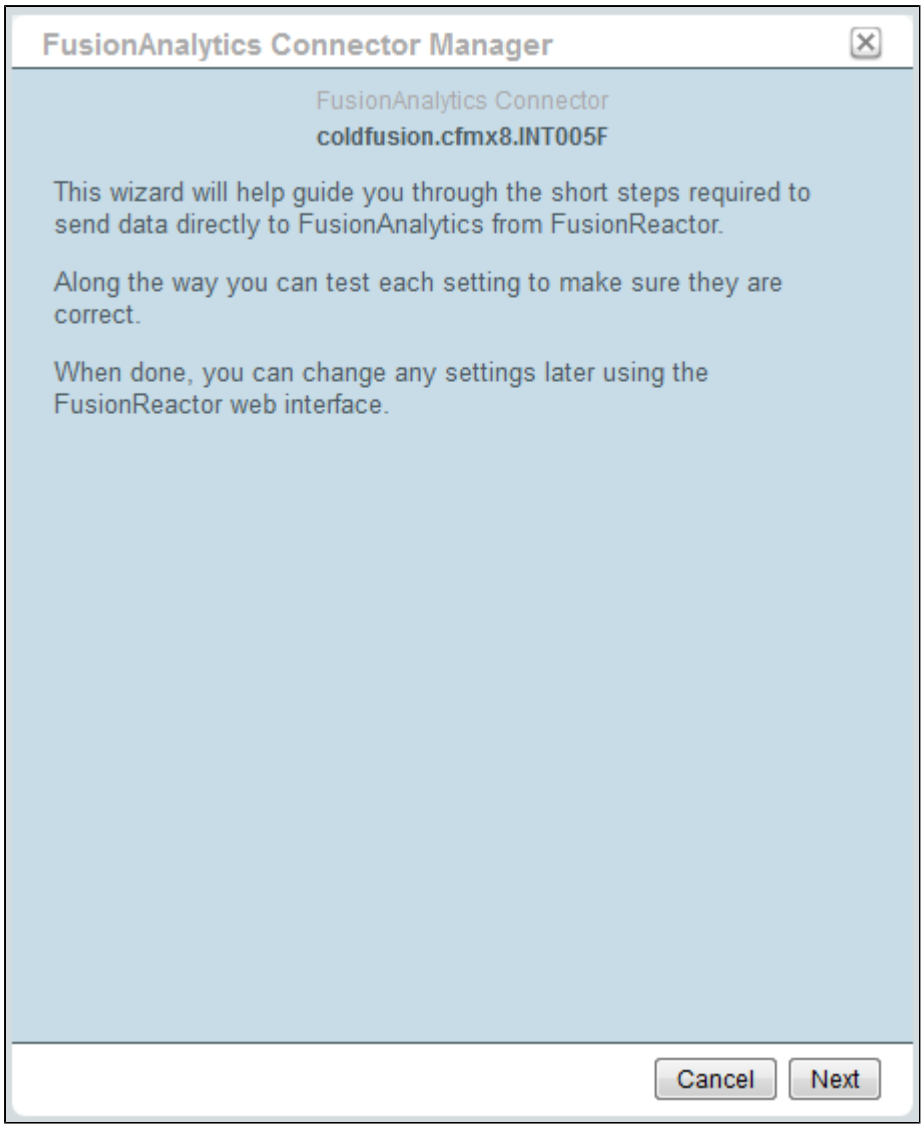


Figure 1: Information

Figure 1: A Welcome screen explaining the FusionAnalytics Connector Wizard is displayed.
Clicking *'Next'* will continue on through the wizard and clicking *'Cancel'* will exit the wizard and take you back to the Instance Manager.

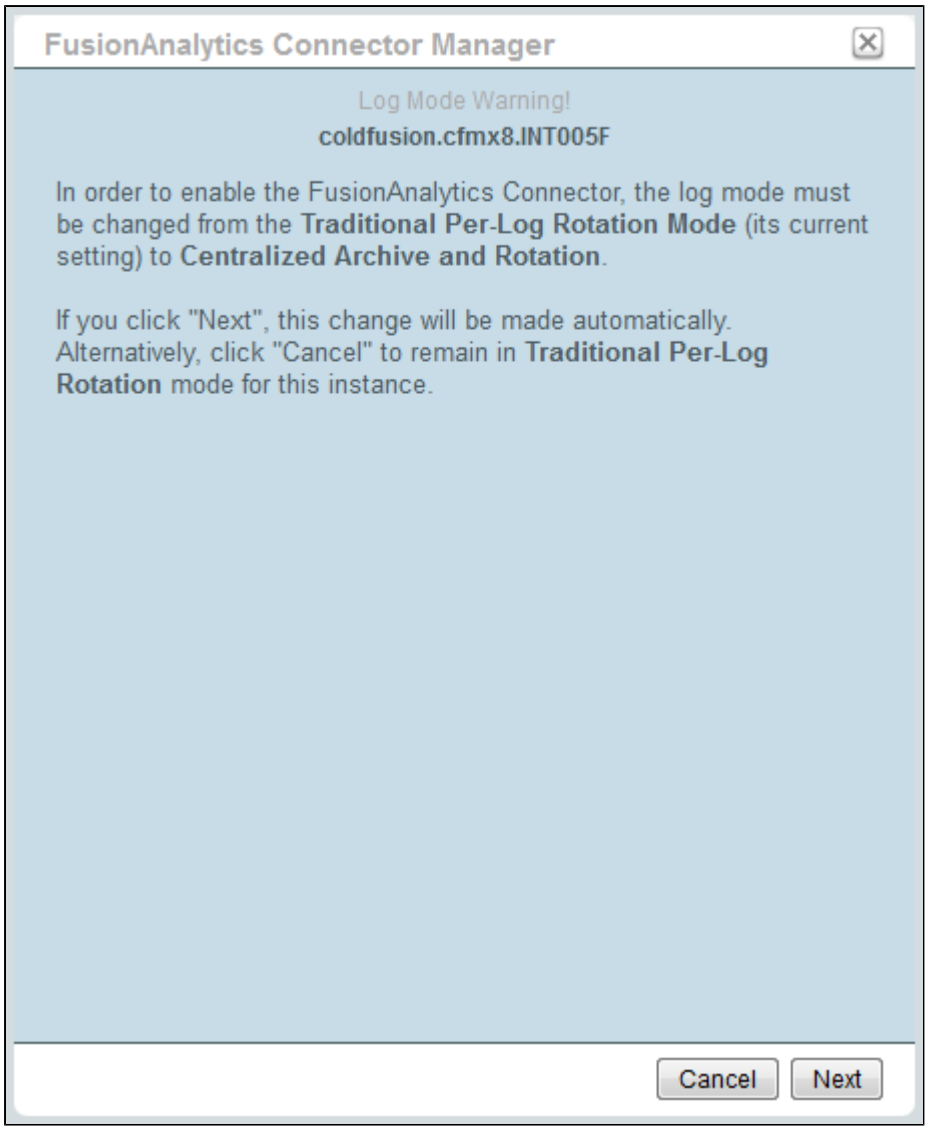


Figure 2: Incorrect log rotation mode

Figure 2: If you have updated a FusionReactor 3.x instance, or you have changed the log rotation mode, this is running in the **Traditional Per-Log Rotation mode.** To enable log file transfer to FusionAnalytics this mode needs to be changed to the new mode introduced in FusionReactor 4.0.0, **Centralized Archive and Rotation mode.** For more information on log rotation modes please see the online documentation covering the Log Settings.

By clicking *'Next'* the log rotation mode will be automatically changed to Centralized Archive and Rotation and the wizard will continue on to Figure 3.

If you do not wish to change the log rotation mode of this FusionReactor Instance at this moment press *'Cancel'* t o return to the Instance Manager.

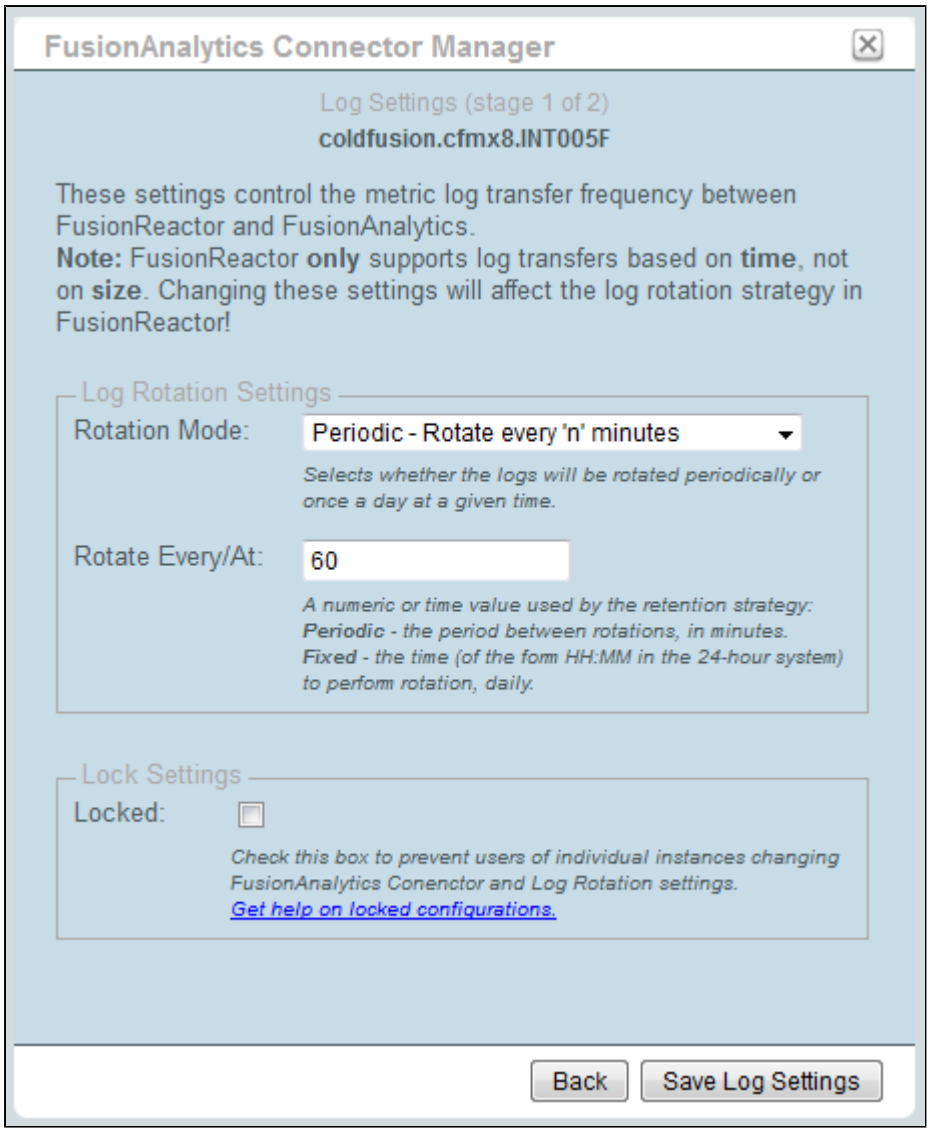


Figure 3: Log Settings

Figure 3: In this screen you are able to configure the Log rotation and Locked settings.

Click *'Save Log Settings'* to save the configured settings and move on to Figure 4.

Click *'Back'* to return to the previous screen, Figure 1.

Log Rotation Settings

Configuration	Description
Rotation Mode	Select whether the logs should be rotated periodically or once a day at a specified time. For more information see the online documentation for Log Settings.

Rotate Every/At	A numeric or time value used by the above Rotation Mode. **Periodic** - the period between rotations, in minutes. **Fixed** - the time (of the form HH:MM in the 24-hour system) to perform rotation, daily. For more information see the online documentation for Log Settings

Lock Settings

Configuration	Description
Locked	Checking this box will prevent users of the FusionReactor Instance modifying vital FusionAnalytics Connector settings. **This means that these settings can only be configured from the FusionReactor Administration Manager.** For more information about locking the FusionAnalytics Connector Settings please see the information about Locked.

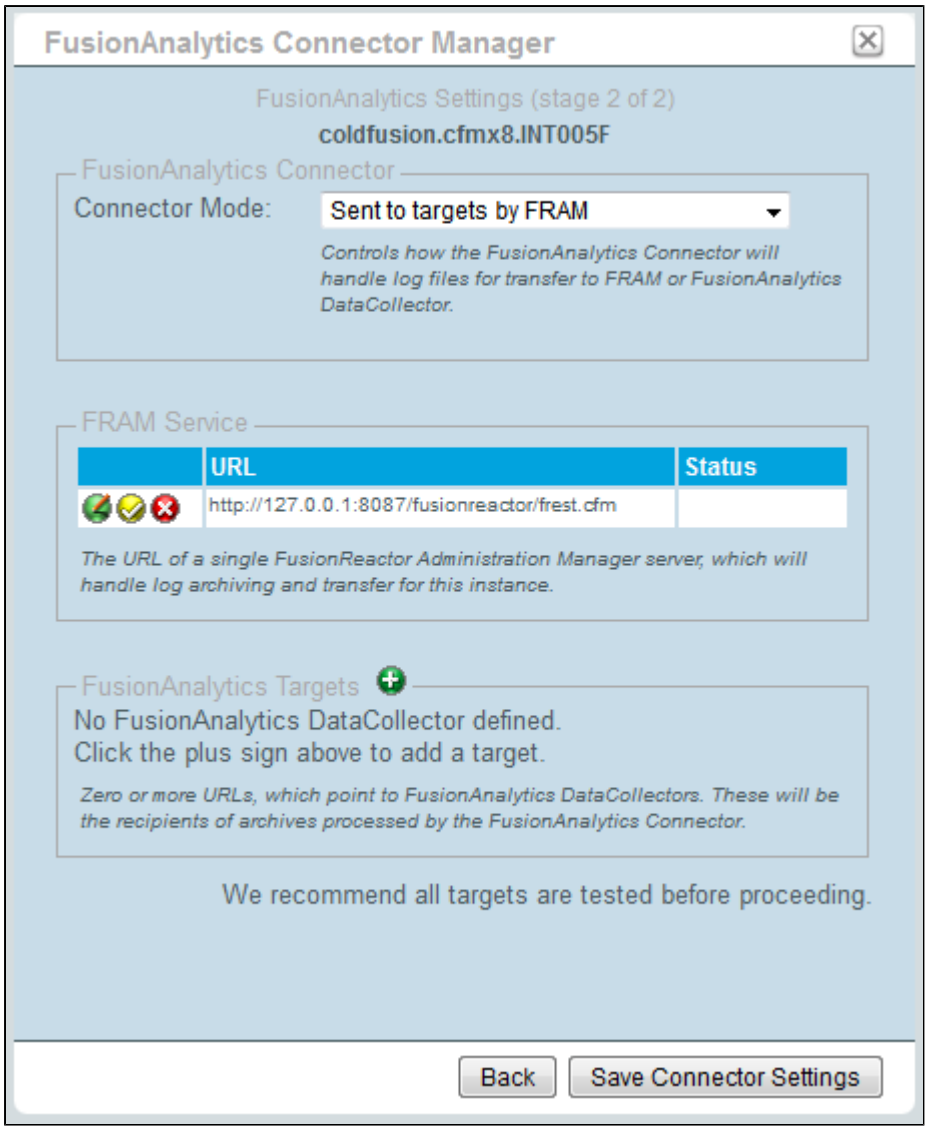


Figure 4: FusionAnalytics Settings

Figure 4: On this screen you have the opportunity to modify the different targets involved in transferring log files to FusionAnalytics, as well as how the log files should be transferred.

Click *'Save Connector Settings'* to save the configured settings and finish the wizard, Figure 13.

Click *'Back'* to return to the previous screen, Figure 3.

FusionAnalytics Connector

Configuration	Description

Connector Mode	This mode controls how the log files are being transferred to the FusionAnalytics DataCollector. **Sent to targets by FRAM** - Log files are compressed and sent to FusionAnalytics targets by FRAM. If FRAM is not available, falls back to **Sent to targets by this instance**. **Sent to targets by this instance** - Log files are compressed and sent locally. **Sent by this instance (uncompressed)** - No compression is performed. Each log file is sent to FRAM or FusionAnalytics directly and uncompressed. **Disabled** - The FusionAnalytics Connector is disabled. No log files are sent to FusionAnalytics. The different modes are described in more details on the page covering the FusionAnalytics Settings

FRAM Service

If the above Connector mode is set to **Sent to targets by FRAM** the FRAM Service is responsible for compressing and sending the log files to the defined FusionAnalytics DataCollector targets. This allows the processing to be offloaded from the instance onto FRAM.

The FRAM Service is by default set to the FusionReactor Administration Manager that created the FusionReactor Instance.

In the above screen you are able to Edit (Figure 5)  , Check (Figure 6)  , Remove (Figure 7)  and Add (Figure 8)  a FusionReactor Administration Manager target.

FusionAnalytics Targets

A FusionAnalytics target is the final receiver of log files. This is a FusionAnalytics DataCollector that processes the generated log files in FusionAnalytics.

To enable log files to be transferred from FusionReactor at least one FusionAnalytics DataCollector Target must be defined.

You are able to Add (Figure 9)  , Edit (Figure 10)  , Check (Figure 11)  and Remove (Figure 12)  a FusionAnalytics DataCollector Target



Figure 5: Edit FRAM target

Figure 5: On this screen you have the opportunity to modify the current FusionReactor Administration Manager (FRAM) target.

Click *'Save FRAM Settings'* to save the configured settings and go back to the FusionAnalytics Settings screen (Figure 4).

Click *'Back'* to return to the previous screen without saving the FRAM target, Figure 4.

Configuration	Description
Host	This is the host running the FusionReactor Administration Manager. By default this is set to the local FRAM Service (127.0.0.1).
Port	This is the port that the FRAM Service is listening on. This is specified in the FusionReactor 4.0.x installation. By default the port for FRAM is 8087.

Username	If FR-REST authentication has been enabled, the FR-REST username needs to be specified. By default FR-REST authentication is enabled locally with no authentication. For more information about FR-REST please see the **FR-REST section** on the Settings page.
Password	If FR-REST authentication has been enabled, the FR-REST password needs to be specified. By default FR-REST authentication is enabled locally with no authentication. For more information about FR-REST please see the **FR-REST section** on the Settings page.
HTTPS	Tick the check box if FRAM is using a secure HTTPS connection instead of regular HTTP. Please see Setting Up FRAM with HTTPS for more information on how to set up FRAM to use a secure connection.
Status	By clicking the button *'Test Target'* the specified FRAM target will be tested. The status of this test will be displayed in the Status column. This can either be **OK** or **Error**. Clicking the status will provide you with more information about the test, such as why the test was unsuccessful.

For more information, please see the online documentation covering FRAM Target.

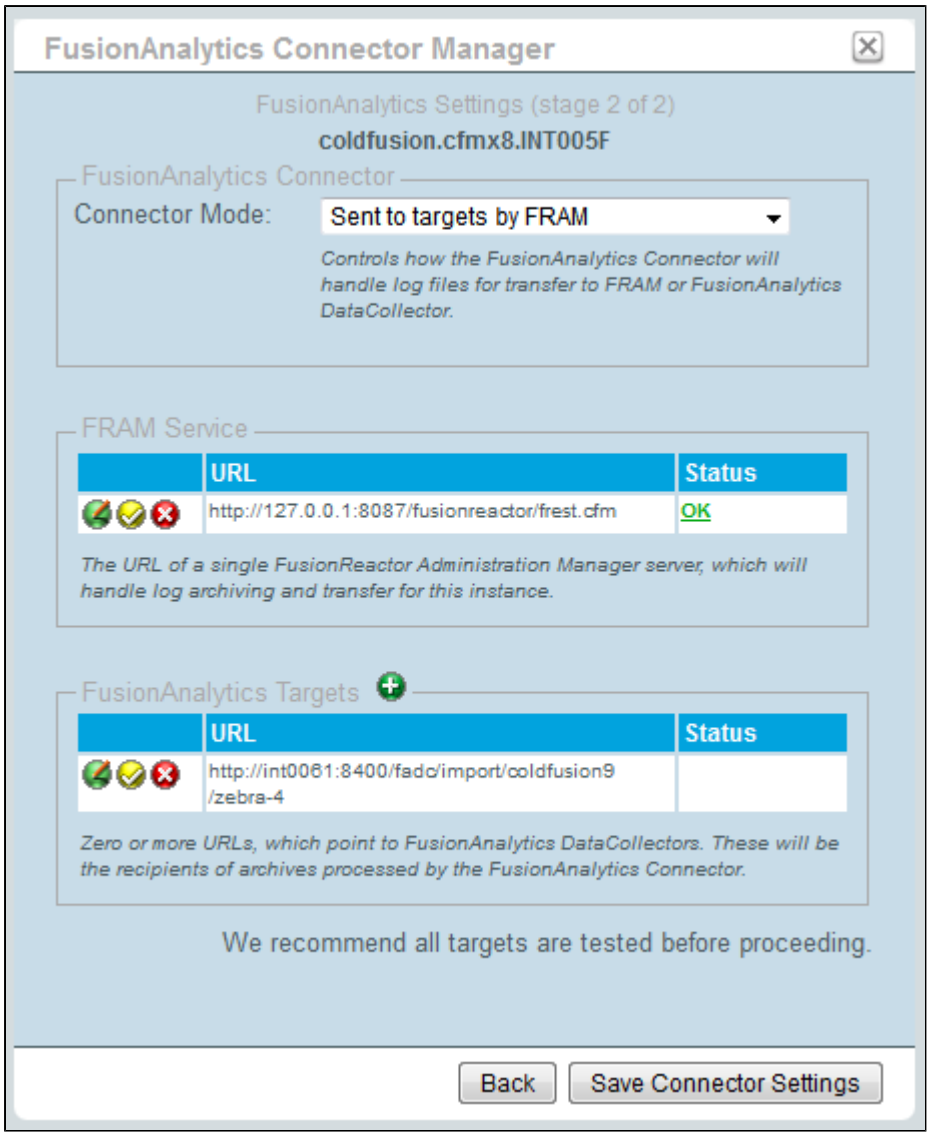


Figure 6: Check FRAM target

Figure 6: After clicking the test button the listed FRAM target will be tested. The status of this test will be displayed in the Status column. This can either be **OK** or **Error**. Clicking the status will provide you with more information about the test, such as why the test was unsuccessful.

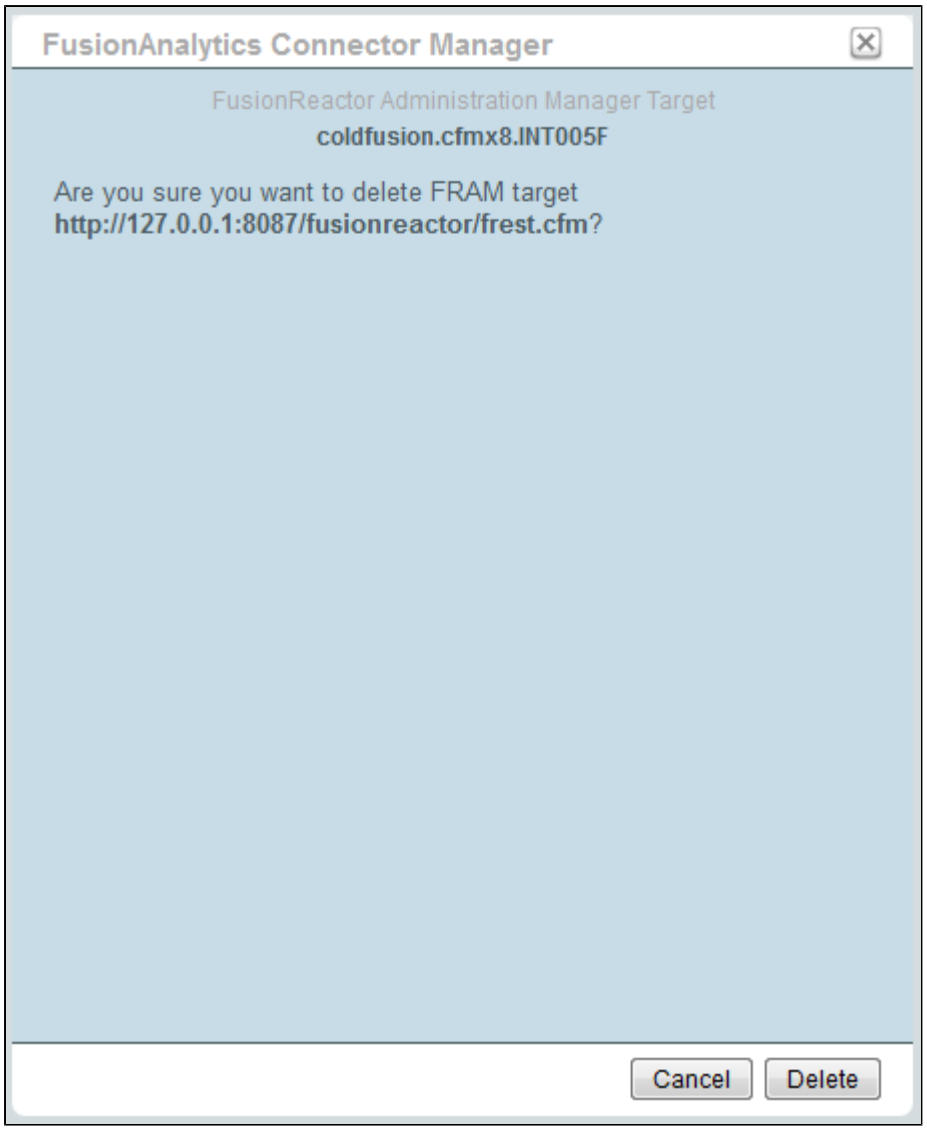


Figure 7: Remove FRAM Target

Figure 7: By clicking for the listed FRAM target you will see a confirmation screen. Click *'Delete'* if you are sure you wish to delete the specified FRAM target. Click *'Cancel'* to go back to the FusionAnalytics Settings screen (Figure 4).



Figure 8: Add FRAM Target

Figure 8: On this screen you have the opportunity to add a new FusionReactor Administration Manager (FRAM) target.

Click *'Save FRAM Settings'* to save the configured settings and go back to the FusionAnalytics Settings screen (Figure 4).

Click *'Back'* to return to the previous screen without saving the FRAM target, Figure 4.

Configuration	Description
Host	This is the host running the FusionReactor Administration Manager. By default this is set to the local FRAM Service (127.0.0.1).
Port	This is the port that the FRAM Service is listening on. This is specified in the FusionReactor 4.0.x installation. By default the port for FRAM is 8087.

Username	If FR-REST authentication has been enabled, the FR-REST username needs to be specified. By default FR-REST authentication is enabled locally with no authentication. For more information about FR-REST please see the **FR-REST section** on the Settings page.
Password	If FR-REST authentication has been enabled, the FR-REST password needs to be specified. By default FR-REST authentication is enabled locally with no authentication. For more information about FR-REST please see the **FR-REST section** on the Settings page.
HTTPS	Tick the check box if FRAM is using a secure HTTPS connection instead of regular HTTP. Please see Setting Up FRAM with HTTPS for more information on how to set up FRAM to use a secure connection.
Status	By clicking the button *'Test Target'* the specified FRAM target will be tested. The status of this test will be displayed in the Status column. This can either be **OK** or **Error**. Clicking the status will provide you with more information about the test, such as why the test was unsuccessful.

For more information, please see the online documentation covering FRAM Target.

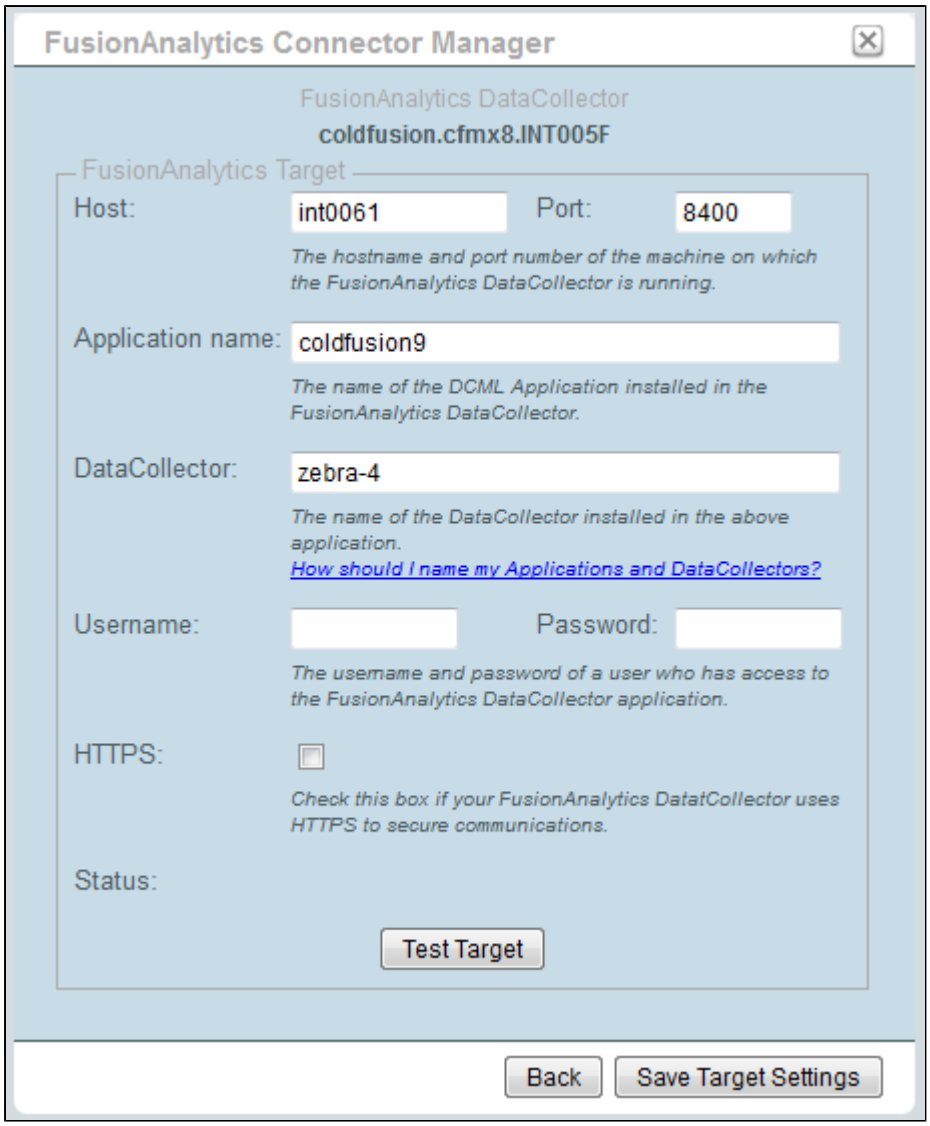


Figure 9: Add a FusionAnalytics Target

Figure 9: On this screen you have the opportunity to add a new FusionAnalytics DataCollector Target.

Click *'Save Target Settings'* to save the configured settings and go back to the FusionAnalytics Settings screen (Figure 4).

Click *'Back'* to return to the previous screen without saving the FusionAnalytics target, Figure 4.

Configuration	Description
Host	This is the host running the FusionAnalytics Server.
Port	This is the port that the FusionAnalytics Server is listening on. This is specified in the FusionAnalytics installation. By default the port for FusionAnalytics is 8400.
Application name	The name of the DataCollector Application that you want to collect the data.

DataCollector	The name of the actual DataCollector that you want to receive the data. Each DataCollector application can have multiple DataCollectors that receive data. The default setting is one DataCollector named faDataCollector-1.
Username	The username for a user who has access to the DataCollector application. This is only used if the DataCollector requires authentication. If you are using the default setting, this is not needed.
Password	The password for the above username.
HTTPS	Tick the check box if the FusionAnalytics DataCollector is using a secure HTTPS connection instead of regular HTTP. Please see Setting Up FusionAnalytics with HTTPS. for more information on how to set up FusionAnalytics to use a secure connection.
Status	By clicking the button *'Test Target'* the specified FusionAnalytics target will be tested. The status of this test will be displayed in the Status column. This can either be **OK** or **Error**. Clicking the status will provide you with more information about the test, such as why the test was unsuccessful.

For more information, please see the online documentation covering FusionAnalytics DataCollector Targets.

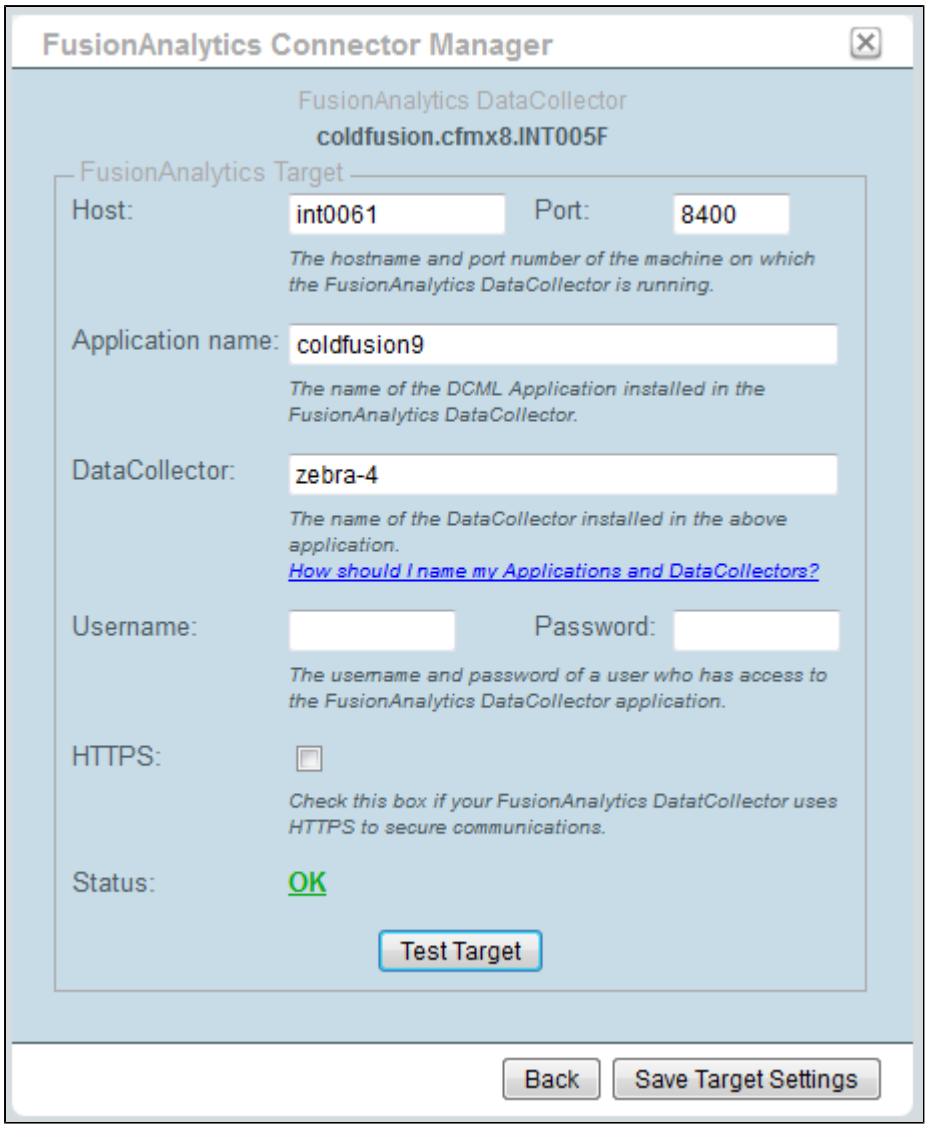


Figure 10: Edit a FusionAnalytics target

Figure 9: On the above screen you have the opportunity to modify the selected FusionAnalytics DataCollector Target.

Click *'Save Target Settings'* to save the configured settings and go back to the FusionAnalytics Settings screen (Figure 4).

Click *'Back'* to return to the previous screen without saving the FusionAnalytics target, Figure 4.

Configuration	Description
Host	This is the host running the FusionAnalytics Server.
Port	This is the port that the FusionAnalytics Server is listening on. This is specified in the FusionAnalytics installation. By default the port for FusionAnalytics is 8400.
Application name	The name of the DataCollector Application that you want to collect the data.

DataCollector	The name of the actual DataCollector that you want to receive the data. Each DataCollector application can have multiple DataCollectors that receive data. The default setting is one DataCollector named faDataCollector-1.
Username	The username for a user who has access to the DataCollector application. This is only used if the DataCollector requires authentication. If you are using the default setting, this is not needed.
Password	The password for the above username.
HTTPS	Tick the check box if the FusionAnalytics DataCollector is using a secure HTTPS connection instead of regular HTTP. Please see Setting Up FusionAnalytics with HTTPS. for more information on how to set up FusionAnalytics to use a secure connection.
Status	By clicking the button *'Test Target'* the specified FusionAnalytics target will be tested. The status of this test will be displayed in the Status column. This can either be **OK** or **Error**. Clicking the status will provide you with more information about the test, such as why the test was unsuccessful.

For more information, please see the online documentation covering FusionAnalytics DataCollector Targets.

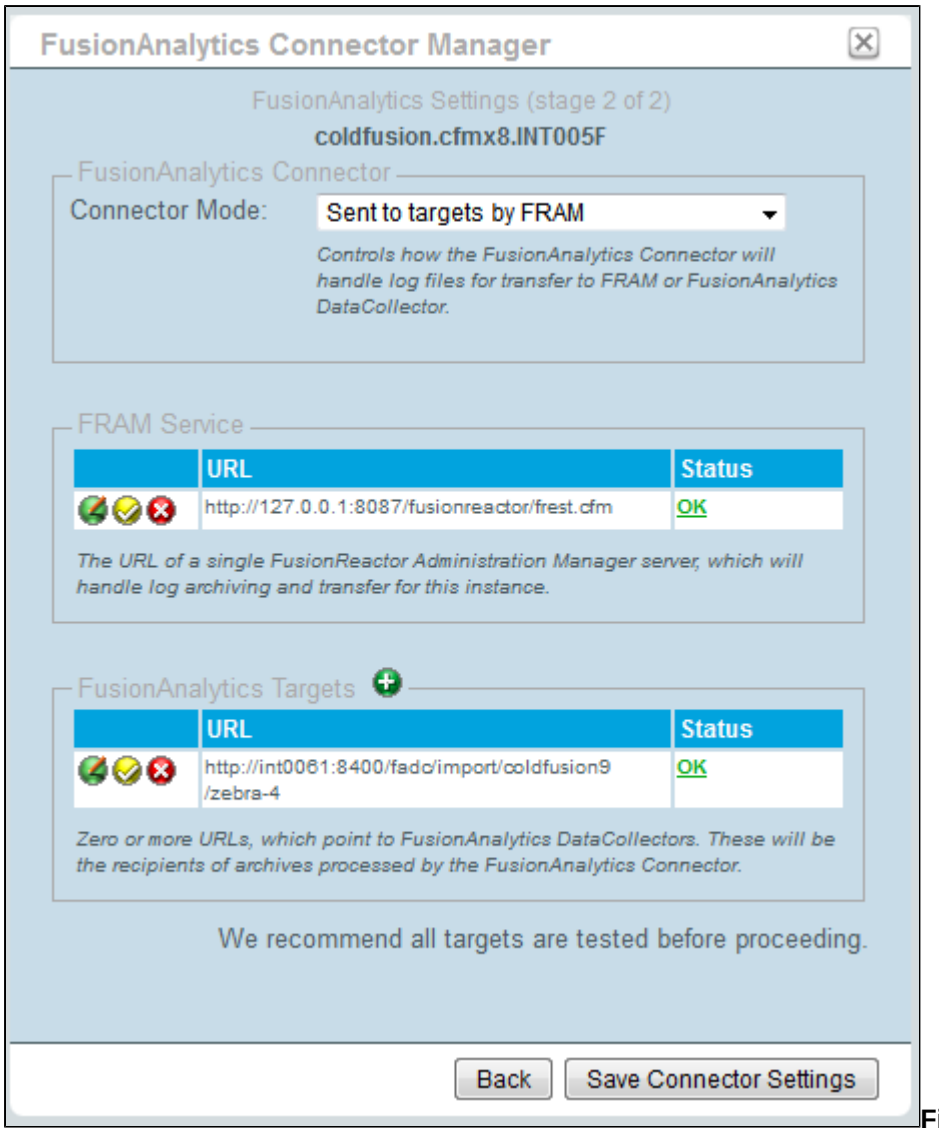


Figure 11: Check a FusionAnalytics Target

After clicking the test button on a listed FusionAnalytics target will be tested. The status of this test will be displayed in the Status column. This can either be **OK** or **Error**. Clicking the status will provide you with more information about the test, such as why the test was unsuccessful.

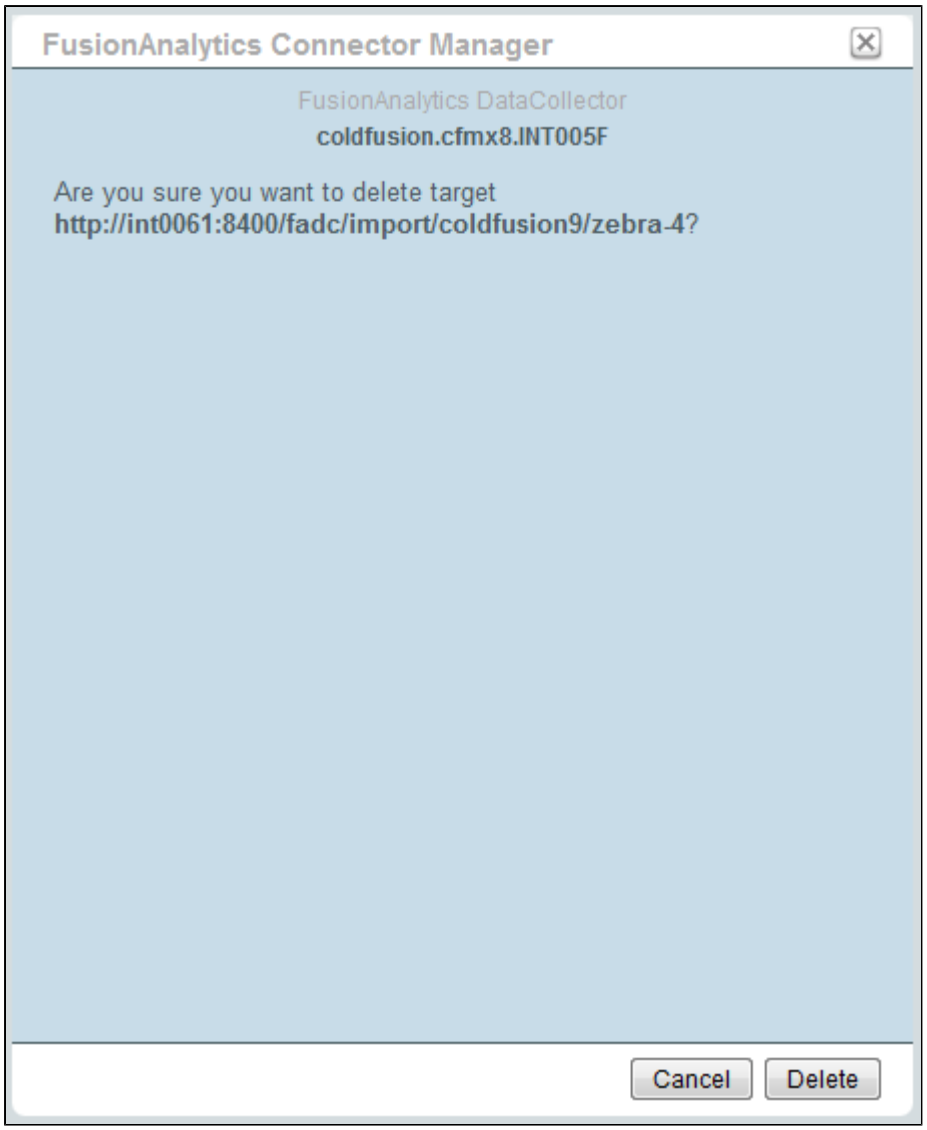


Figure 12: Removing a FusionAnalytics Target

Figure 12: By clicking the remove for a listed FusionAnalytics target you will see a confirmation screen. Click *'Delete'* if you are sure you wish to delete the specified FusionAnalytics target. Click *'Cancel'* to go back to the FusionAnalytics Settings sreen (Figure 4).

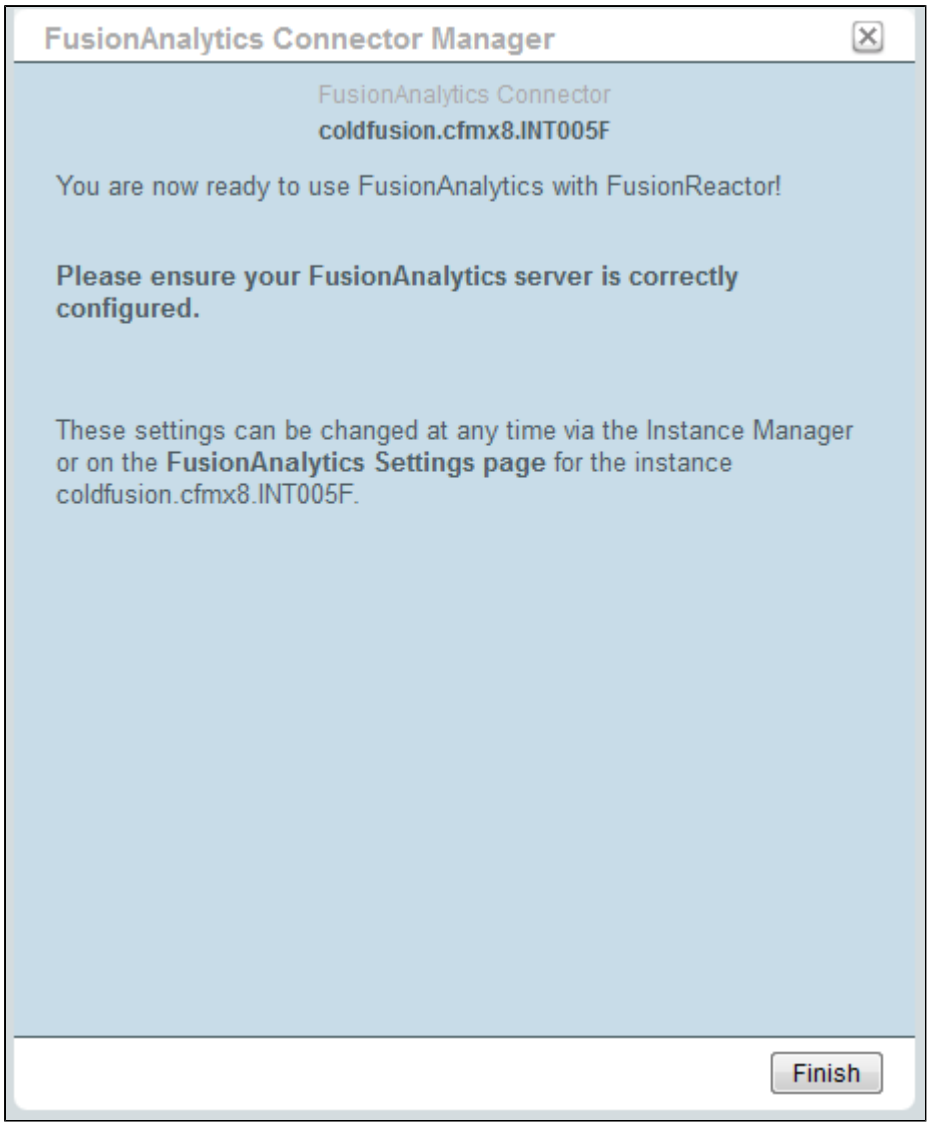


Figure 13: Configuration is complete.

Figure 13: The FusionAnalytics Connector has been configured for the specified FusionReactor Instance.

> Make sure that the configured FusionAnalytics Server is running and correctly configured before transferring logs.

Click *'Finish'* to return to the Instance Manager.

Update FusionReactor Instance

Inside the Instance Manager, you will be able to update any FusionReactor Instance if it has a previous version of FusionReactor installed. To do this, click on the button next to the desired FusionReactor instance or click the button in the rightmost column header (You need to tick the check boxes of the desired FusionReactor Instances or click the hyperlink *"Select All"* if you use the button in the column header). You can also click the *"U pdate"* hyperlink next to the version number of an instance that can be updated.

> **If FusionReactor was installed on a server as a split-jar install (according to Technote FRS-240) it needs to be correctly updated.**
>
> Please follow the guide described in the following Technote (FRS-258) to successfully update FusionReactor

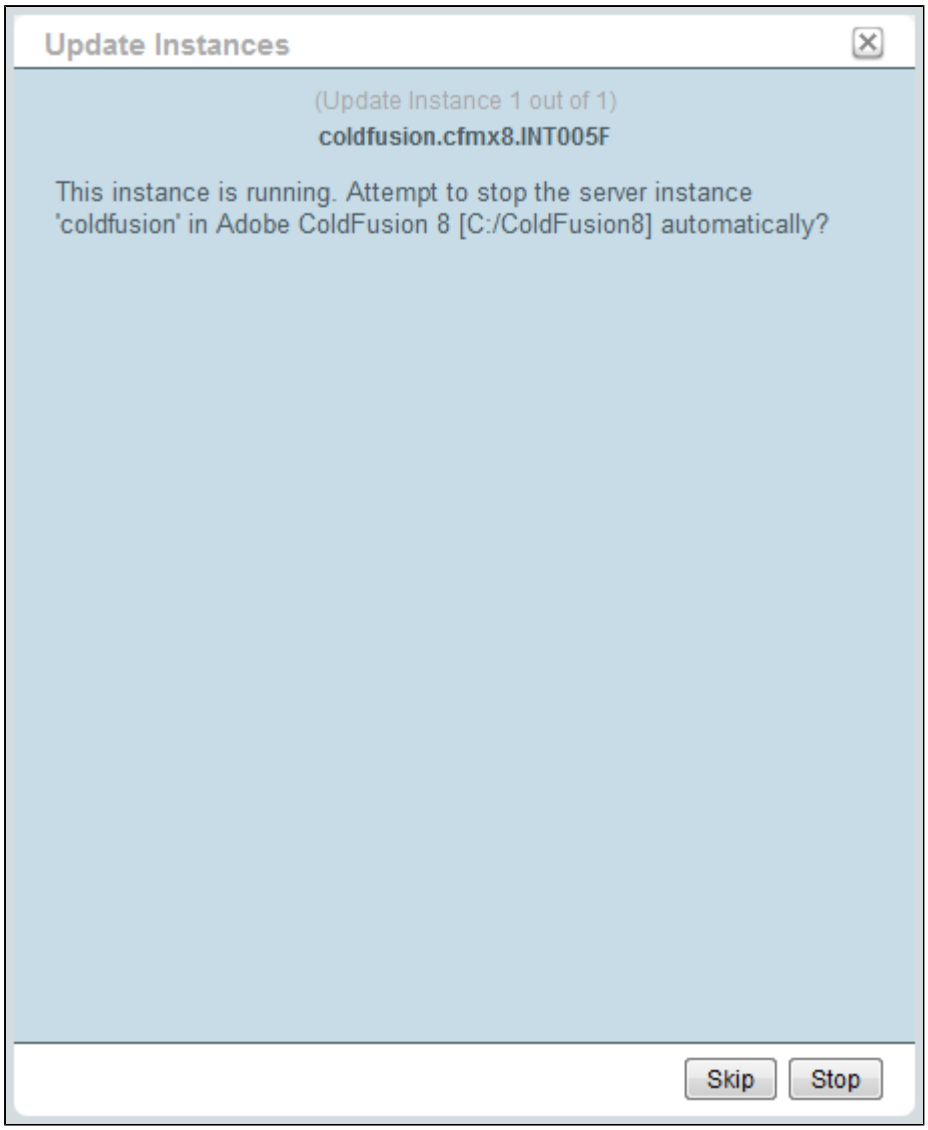


Figure 1: Server Instance Running

Figure 1: If the server instance you are attempting to update FusionReactor on is currently running, it needs to be stopped in order to update the FusionReactor instance.

By clicking *'Stop'* the Instance Manager can attempt to stop the server instance automatically (See Figure 2).

If you wish to stop the server at another point, or wish to update it at a later point, click *'Skip'*. If you have selected to update more than one FusionReactor instance, the wizard will advance to the next FusionReactor instance. Otherwise the Finish screen (Figure 6) will be displayed.

> To update the FusionReactor Instance the server instance must be completely shut down.

If the server instance you are updating FusionReactor on is not running Figure 3 will be displayed.

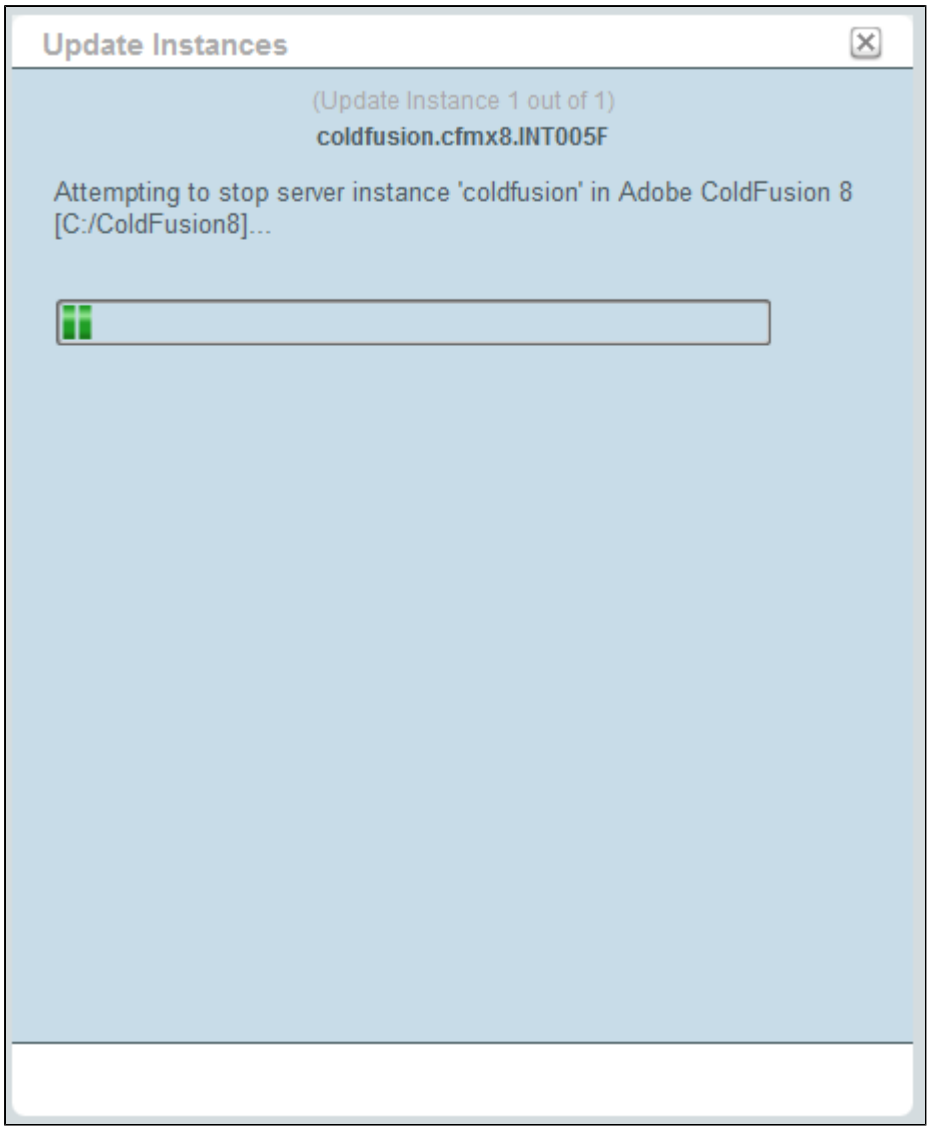


Figure 2: Server Instance Stopping

Figure 2: FusionReactor is attempting to stop the server instance. Please note that stopping the server may take several minutes depending on your system and server.

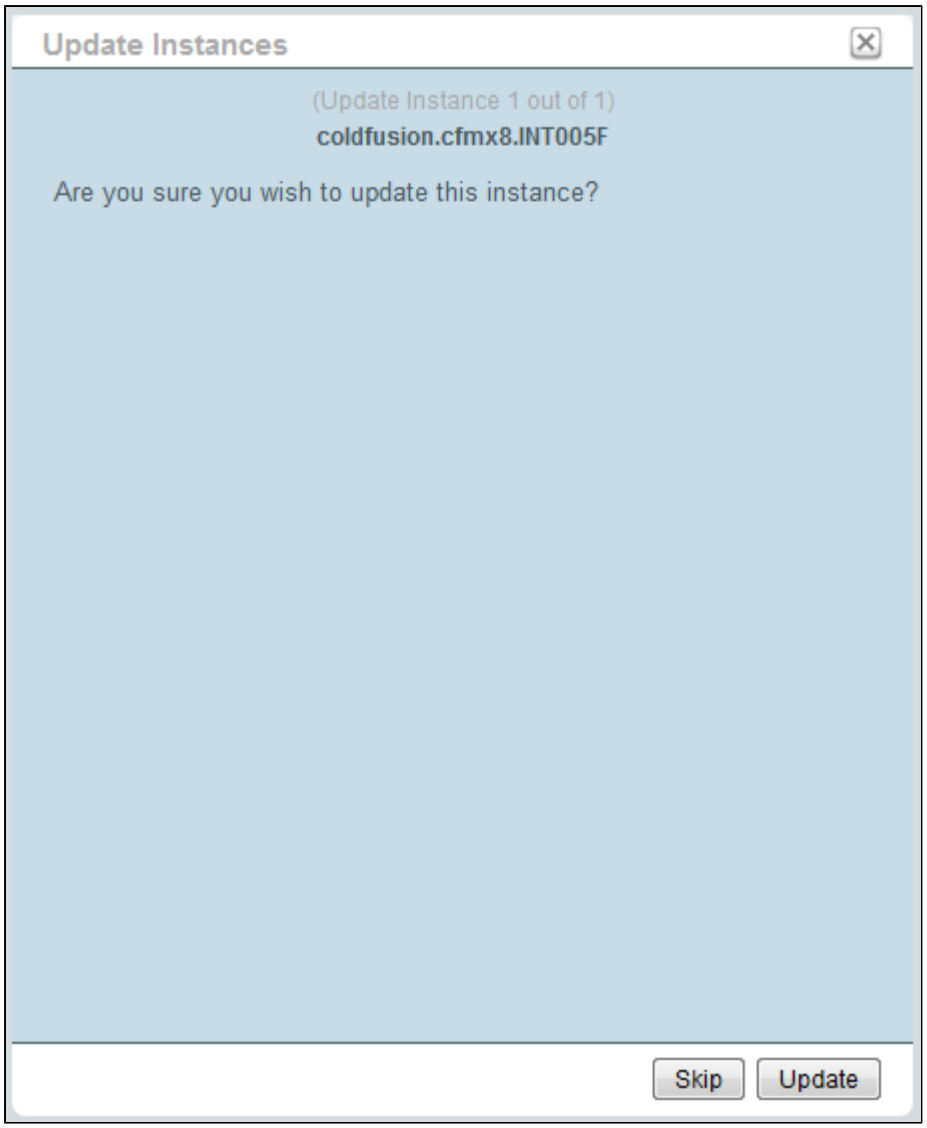


Figure 3: Update Confirmation

Figure 3: By clicking *'Update*' the FusionReactor Instance will be updated to the same version of FusionReactor as is running on the FusionReactor Administration Manager (FRAM).

If you do not wish to update the FusionReactor Instance at this point press *'Skip'*. If you have selected to update more than one FusionReactor Instance the wizard will advance to next FusionReactor instance, if not the finish screen will be displayed (See Figure 7).

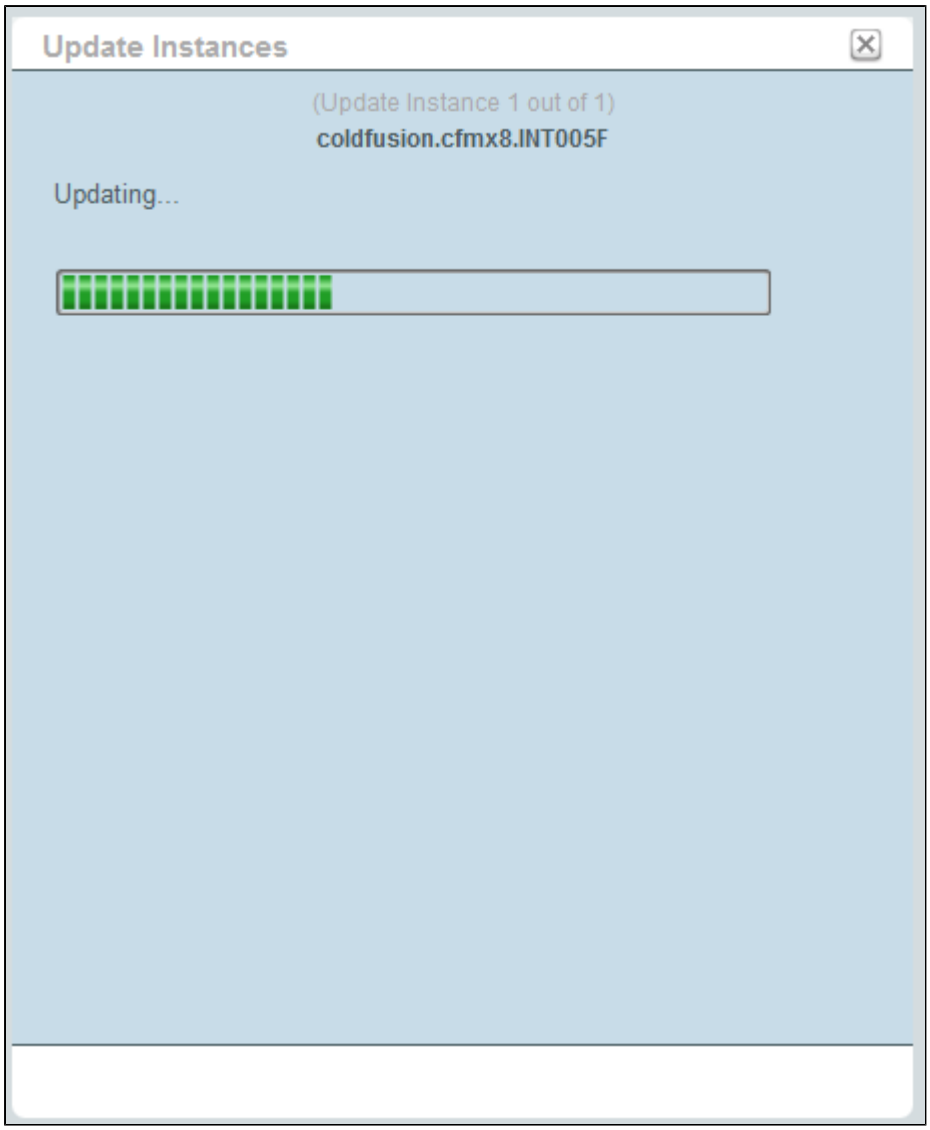


Figure 4: Update Progress

Figure 4: FusionReactor is updating the selected FusionReactor Instance. Updating the FusionReactor Instance should only take a couple of seconds, please note that the time may vary depending on your system.

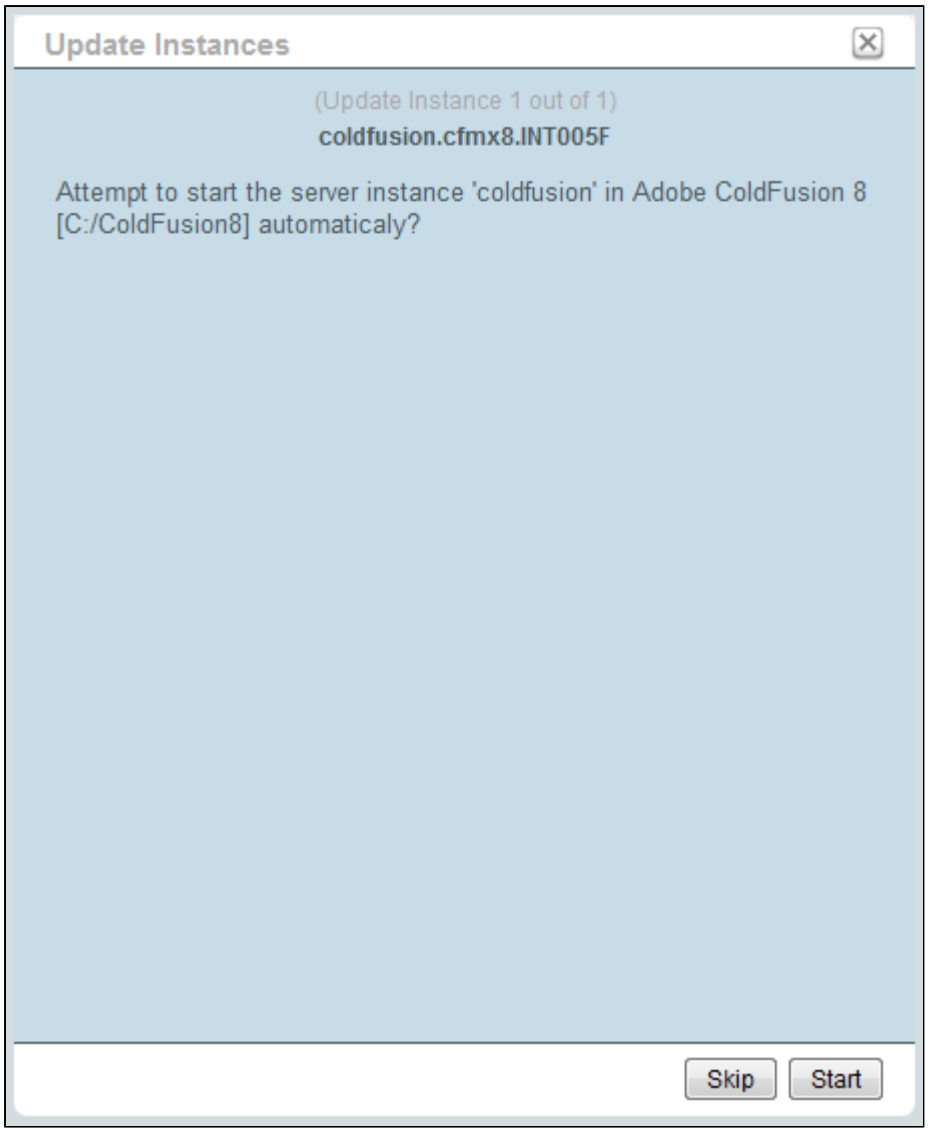


Figure 5: Server Start

Figure 5: By Clicking *'Start'* the Instance Manager can attempt to automatically start the server instance again (See Figure 6).

If you wish to start the server at another point, or wish to manually start it, click *'Skip'* (See Figure 7).

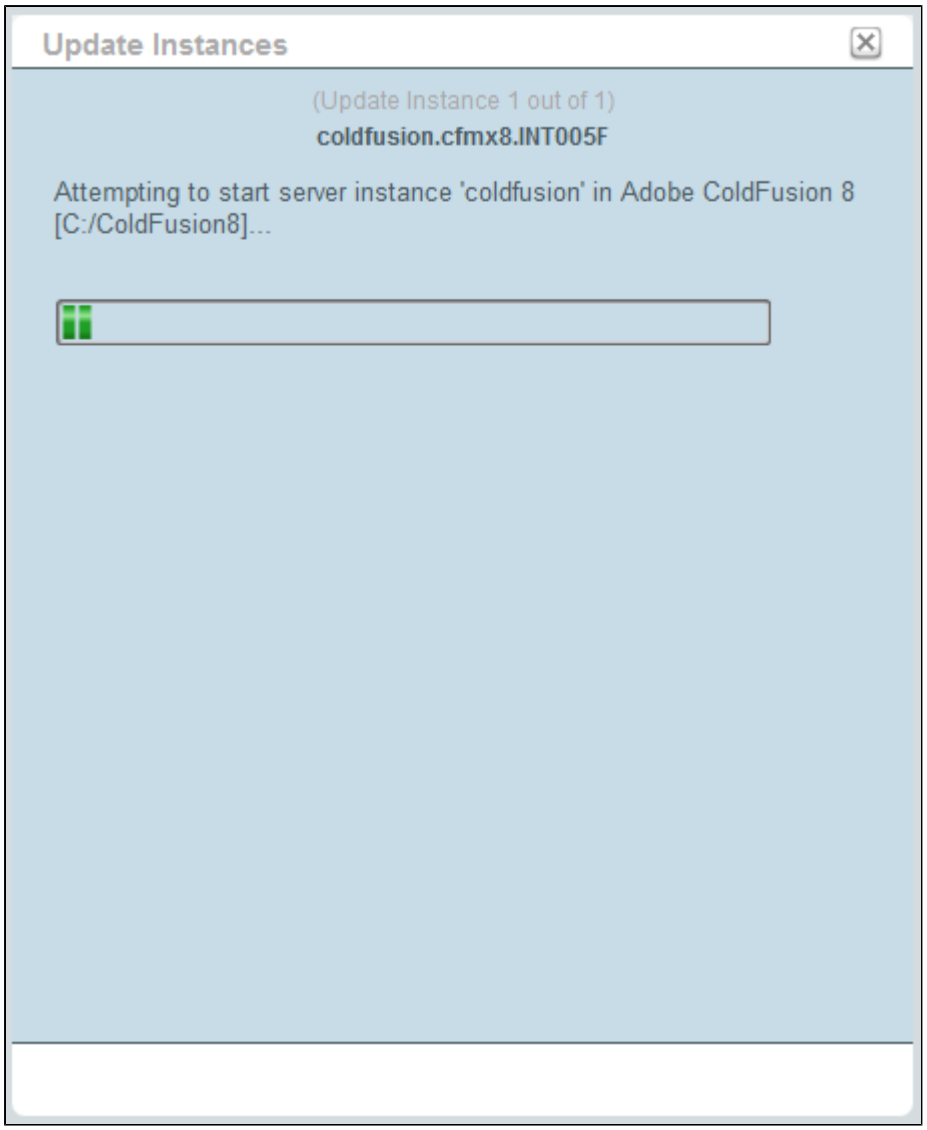


Figure 6: Server Instance Starting

Figure 6: FusionReactor is attempting to start the server instance. Please note that starting the server may take several minutes depending on your system and server.

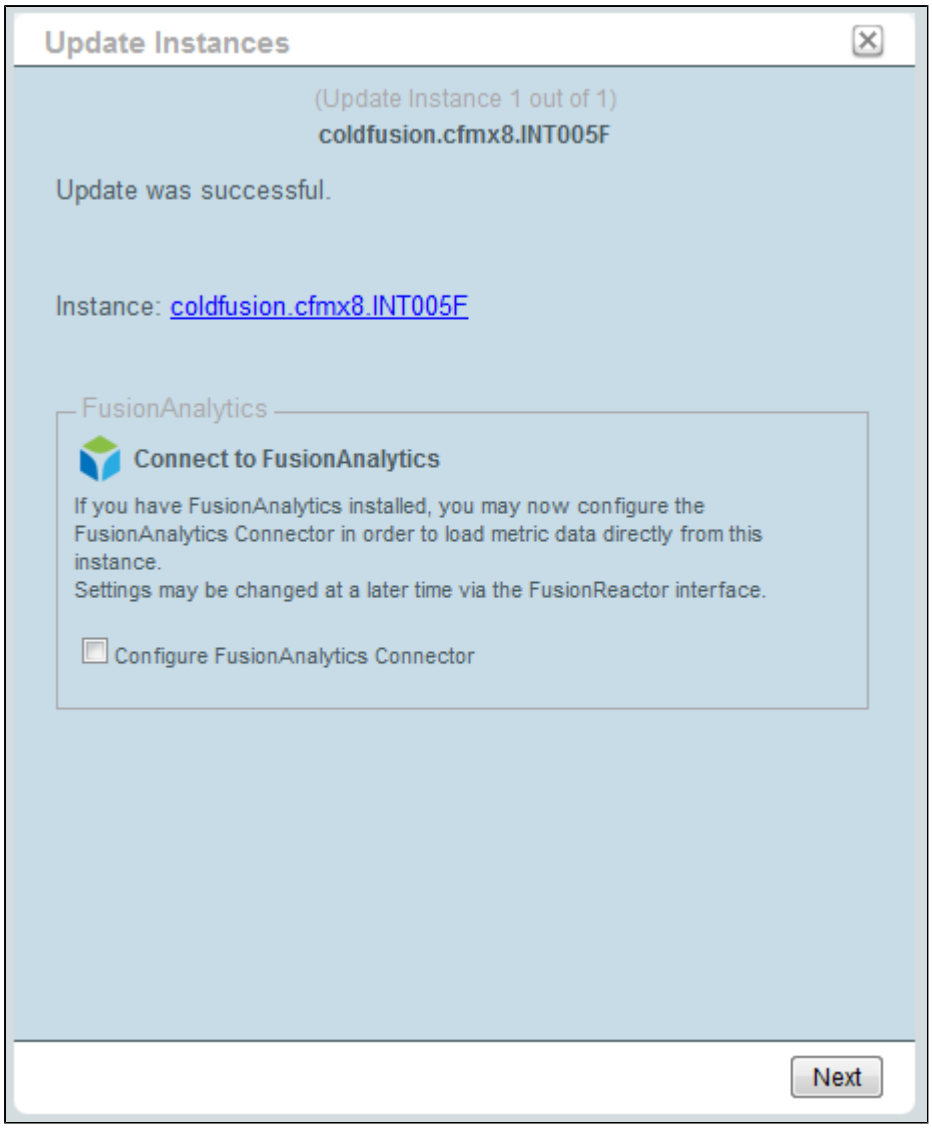


Figure 7: Update Status

Figure 7: Once the update is complete you will have the choice to configure the FusionAnalytics Connector. Tick the check box *'Configure FusionAnalytics Connector'* and click *'Next'* if you wish to configure log transfer to FusionAnalytics now. Otherwise leave the check box unticked.

If you have selected to update multiple FusionReactor instances clicking *'Next'* will advance the wizard to install the next instance.(Go back to Figure 1).

If only one instance was selected you have finished the Installation wizard as displayed by Figure 9.

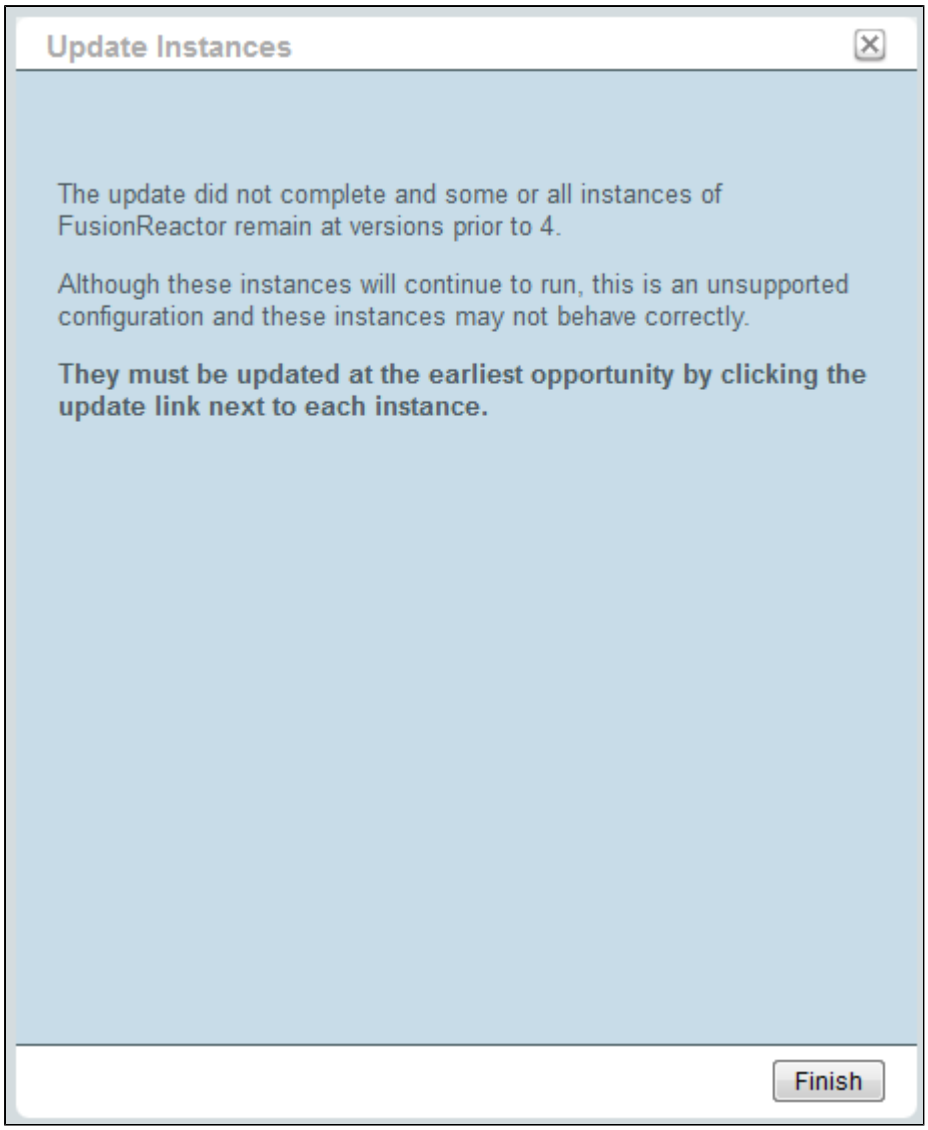


Figure 8: Finished Updates

Figure 8: If the following screen is displayed you have instances that are still on aversion prior to FusionReactor 4.

> Any instances running a version of FusionReactor < 4.0.0 must be updated at the earliest opportunity as this is an unsupported configuration. Please see the following page for more information on Unsupported Configurations.

Click *'Finish'* to return to the Instance Manager.

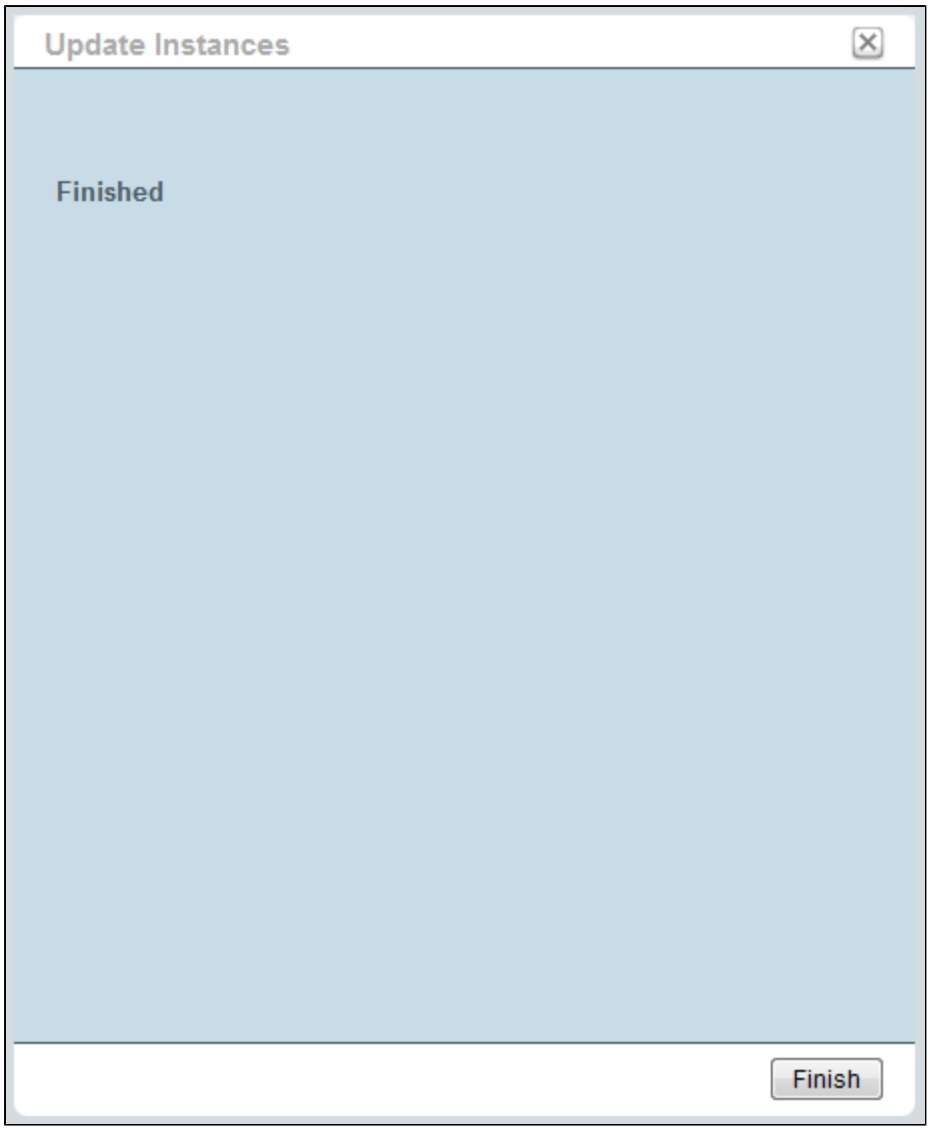


Figure 9: FusionReactor update finished

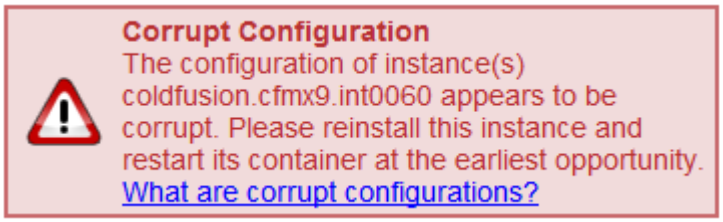
Figure 9: The selected FusionReactor Instances have been updated. Click *'Finish'* to return to the Instance Manager.

Corrupt Configurations

Identifying Corrupt Configurations

FusionReactor's **Instance Manager** (available in the *FusionReactor Administration Manager*) sometimes needs to manipulate the configurations of the individual instances it's monitoring.

To make this operation as safe as possible, some checks are done prior to any edit operation. If any of these checks fail, FRAM refuses to perform the edits and displays a warning, like the following:

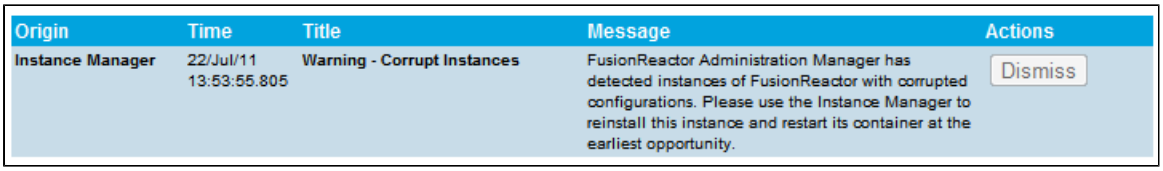


Because this is an exceptional, potentially unsafe situation, you'll also get a non-dismissable notification, flagged

in the table contents:



Clicking on this notification will take you to a page which details the error:



Origin	Time	Title	Message	Actions
Instance Manager	22/Jul/11 13:53:55.805	Warning - Corrupt Instances	FusionReactor Administration Manager has detected instances of FusionReactor with corrupted configurations. Please use the Instance Manager to reinstall this instance and restart its container at the earliest opportunity.	Dismiss

Corrupt configurations can be detected for a number of reasons:

- Completely missing configuration data (`reactor.conf` and `fram.conf` files)
  - Or the files are empty.
- The version of FusionReactor recorded in the configuration is different to that of FRAM (FRAM must be used to update these instances to the new current version of FusionReactor).
- Certain keys required for the operation of FRAM are missing.

Remedying Corrupt Instances

The most failsafe way to remedy corrupt instances is to have **Instance Manager** remove them, then add them again fresh. The J2EE container should be restarted as part of this process.
You'll lose any configuration you've made up to that point, but the configuration will once again be usable by FRAM.

Also See

Uninstall FusionReactor Instance

Unsupported Configurations

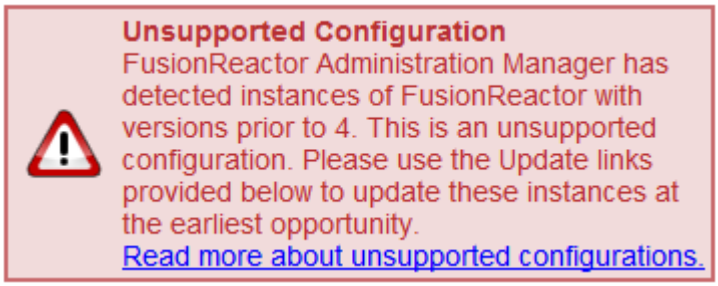
What are Unsupported Configurations?

FusionReactor 4 has to impose some constraints on the overall configuration of instances within a FRAM-administered server. This was an unavoidable limitation and is a side-effect of the highly diverse nature of FusionReactor installations.

An unsupported configuration is reported in the table of contents with a system notification:



... and in the **Instance Manager**, with a warning box:

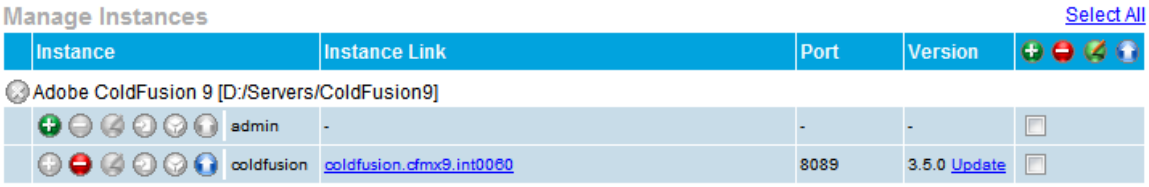


In this case, FusionReactor Administration Manager is reporting that it has detected a version of FusionReactor prior to version 4.

Prior Versions are Unsupported

The structure of FusionReactor has been refined in version 4. More functionality has been introduced, and the instance directory structure has been extended.

For this reason, the installation and initial update process tries very hard to update prior 3.x version of FusionReactor to 4.0.0. If, for whatever reason, this process failed, FusionReactor will warn you by displaying this dialogue on the Instance Manager. The offending instances will be flagged and will have a blue "Update" button, together with a link in the table:



It's important to note that your original 3.x instances might continue to run, but they should be updated as soon as possible using the link provided. We don't support a mixed configuration of 4.x and pre-4.x instances.

The system notification associated with this condition is non-dismissable, but will be removed if the affected instances are updated.

See Also

Update FusionReactor instance

Setting Up FRAM with HTTPS

FusionReactor Administration Manager (FRAM) supports secure (HTTPS) connections, however you need to setup / configure the server to use HTTPS for this to work.

Tomcat

You must install a Certificate into Tomcat to enable HTTPS. HTTPS then runs on a different port to the HTTP port (typically 8087)

- Modify the `server.xml` file located under `[FusionReactor]/tomcat/conf`

```
<Connector port="8443" protocol="HTTP/1.1"
SSLEnabled="true"
maxThreads="150" scheme="https"
secure="true"
keystoreFile="C:/keys/keystore.jks"
keystorePass="changeit"
clientAuth="false" sslProtocol="TLS" />
```

- This is using port 8443 and assuming the keystore password is "changeit". Make any modifications as necessary
- Restart the server (start it using catalina.bat on Windows, catalina.sh on Linux to make sure no errors are being displayed).
- For more information on how to configure a tomcat server to use SSL see: http://tomcat.apache.org/tomcat-6.0-doc/ssl-howto.html

Creating a Self-Signed Certificate

- You need to have Java installed and access to the keytool (java/bin)
- Run the following statement from the command line:
  `keytool -genkey -keyalg RSA -alias selfsigned -keystore keystore.jks -storepass changeit -validity 360 -keysize 2048`
- When prompted, enter "localhost" as the first and last name!
- Use the generated keystore in your tomcat installation.

Accessing FusionReactor Administration Manager (FRAM)

If you have followed the above guide, FRAM should now be available on the HTTPS protocol, accessible on port 8443.
You should be able to access FRAM on a URL similar to: https://127.0.0.1:8443/fusionreactor/

Licensing

Uploading your License File

Activation Period

License Activation

Activating Multi-Key License Files

License Policy

Uploading your License File

To upload a license file, go to the About page under FusionReactor in the Table of Contents within either FRAM or an instance of FusionReactor and you will see a form with a file upload button (Figure 1).



Figure 1: License Upload

Figure 1: Click the '*Browse...*' button and locate the license file. The license has the file extension '*.lic*', for example it may be called `FR3Y1ENT-FR-0013284.lic`.

Once entered, click the '*Upload New License File*' button to upload to the license to FusionReactor.



Figure 2: License Upload

Figure 2: After submitting you will need to login again. You will see your new license details on the middle of the screen straight away when you log in and you will be able to find the same details on the **About** page under FusionReactor in the Table of Contents.You will then be given 10 days to activate the license. If the computer with FusionReactor is connected to the internet you can click on the "Activate Current License" link to activate it. For information on activating a license whilst you are either online or offline please see License Activation to find out how to activate your license.

Activation Period

Once you have uploaded a license key you have 10 days to activate that license. It is recommended that you activate your license immediately after upload. If you do not activate your license within 10 days you will be unable to use FusionReactor until you activate the license. The next time you log into FRAM, you will be shown the image below (Figure 1).

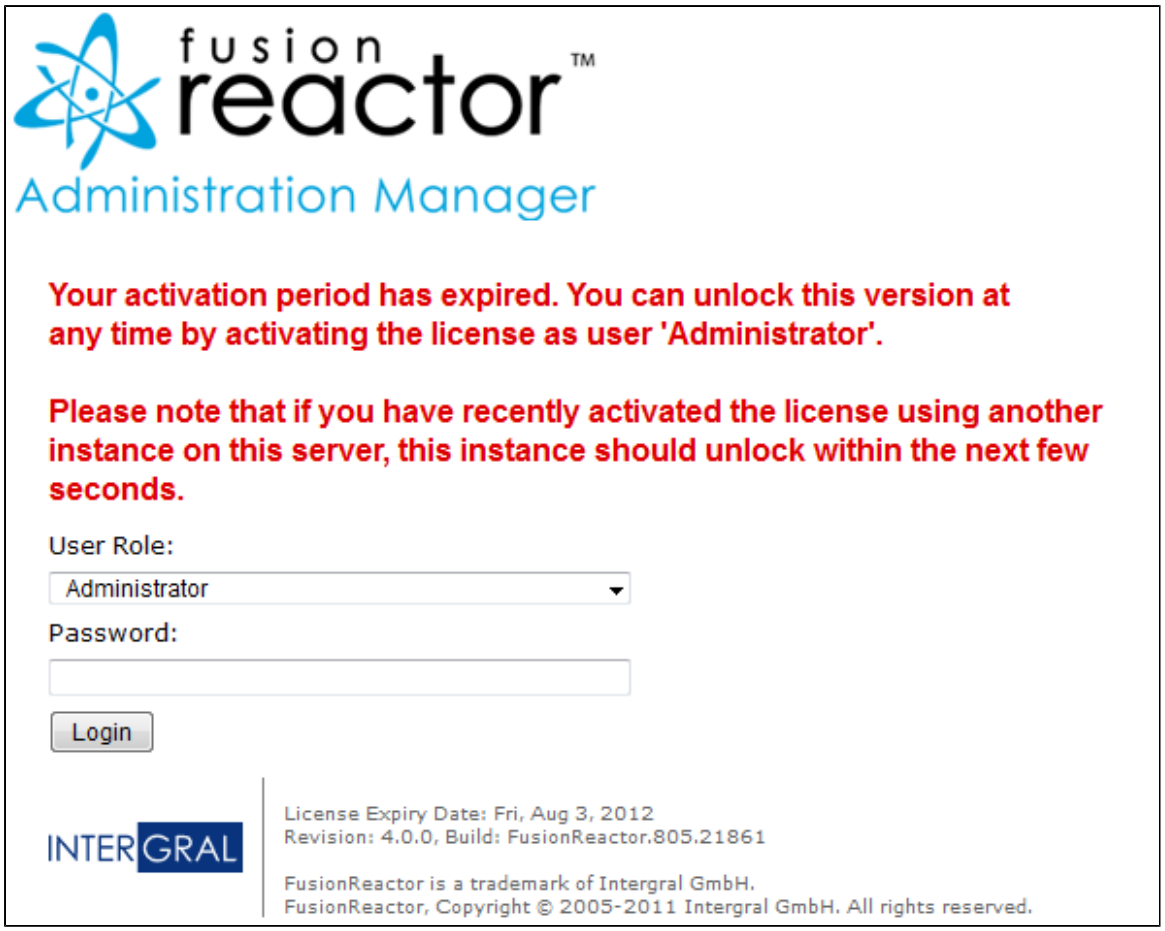


Figure 1: Activation Expired

To activate the license after the 10 days activation period has expired you must login as the **Administrator** and then click on "Activate current license". If your license has ran out of activations, you will be shown Figure 2. Otherwise you will be shown Figure 3.

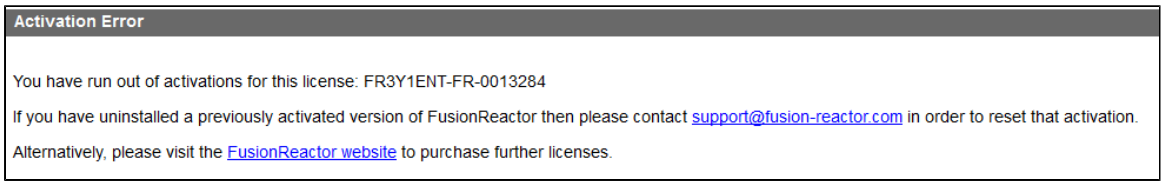


Figure 2: Activation Error

Figure 2: If you have ran out of activations, you will no longer be able to use that particular license unless you've previously uninstall an activated version of FusionReactor. You will be able to contact support@fusion-reactor.com to have that activation reset. Otherwise, you will need to purchase a new license which you can do at the FusionReactor website.

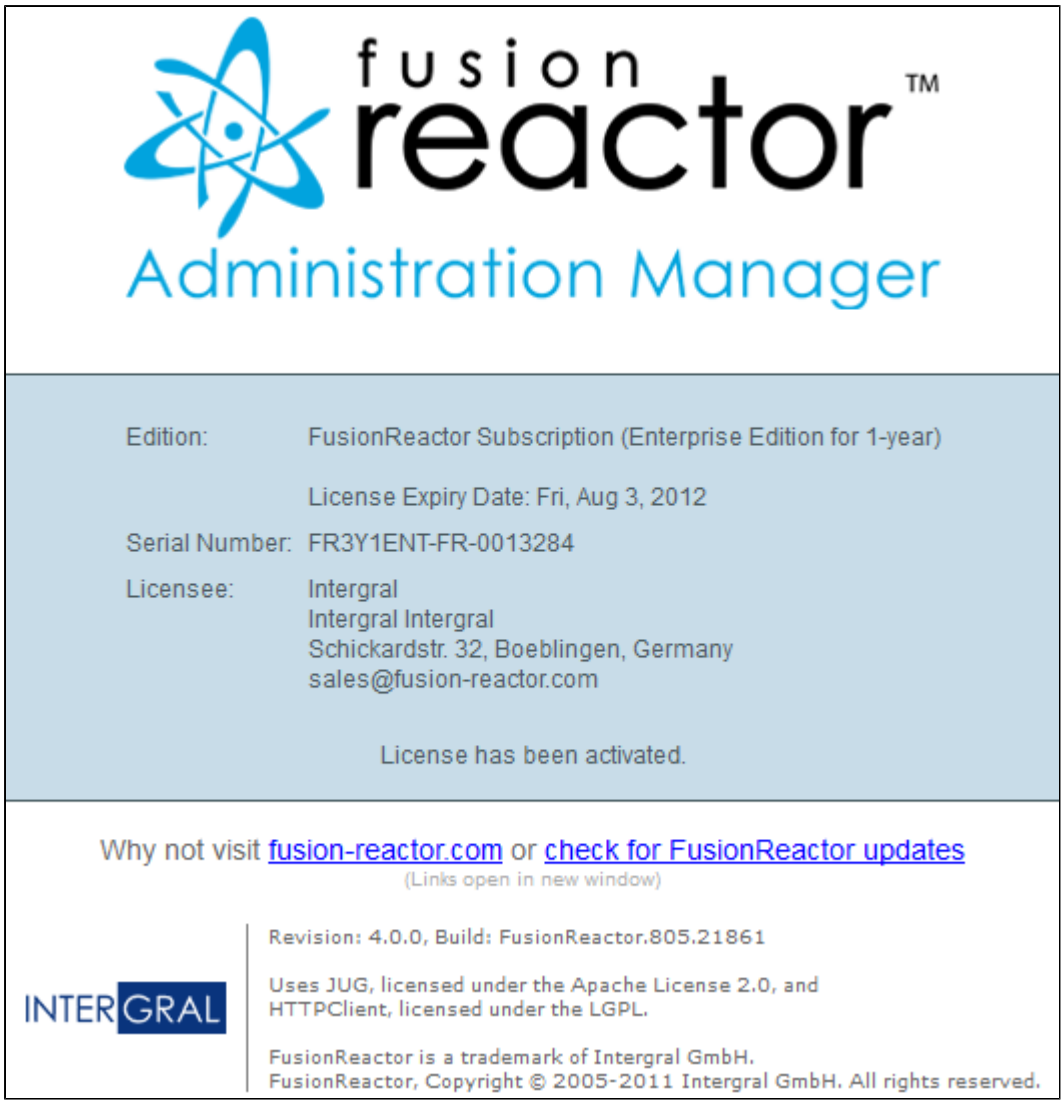


Figure 3: Successful License Activation

Figure 3: You have successfully activated your license.

License Activation

License Activation - Automatic

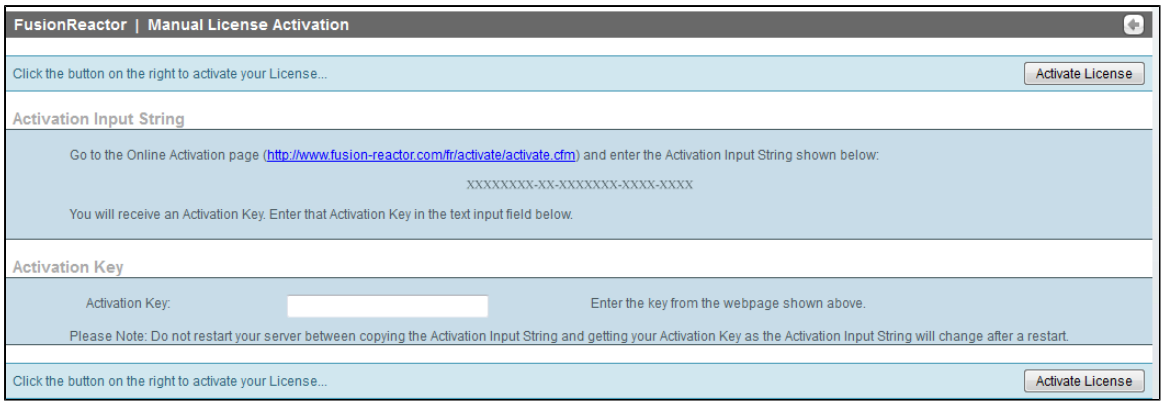
You can activate in FRAM by clicking on FusionReactor in the table of contents and then clicking the page titled About. Once on that page you will see a link titled *'Activate current license'*, clicking this will activate your license.

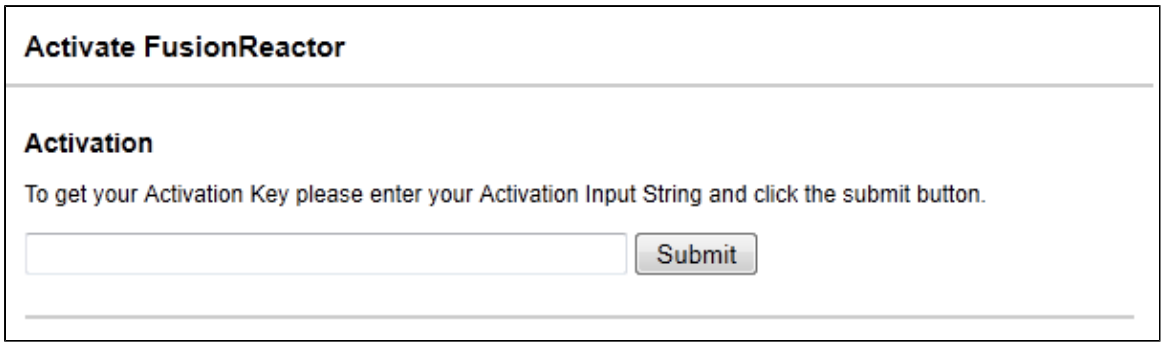


License Activation - Manual

If you do not have a direct internet connection on the computer you would like to install FusionReactor on, you will have to manually activate your license. Please follow the instructions below to manually activate your license.

1. Download FusionReactor on a computer with internet access and copy it to your computer without internet access to install.
2. FusionReactor will recognize you are not connected to the internet and the screen below will appear.
3. Make a note of the "Activation Input String" and then move over to a computer with internet access.
4. Click on the link above the "Activation Input String" (http://www.fusion-reactor.com/fr/activate/activate.cfm) , this will take you to an activation page.
5. Input the "Activation Input String" into the Activation text field and press the Submit button, you will then be given an activation key. Make a note of this.
6. Return to the computer with FusionReactor installed and copy out the activation key into the text field titled "Activation Key".
7. Click Activate License to activate your license.





This is where you input your *'Activation Input String'*.

See Also

About

Activating Multi-Key License Files

A single license file may contain multiple license keys, one for each license you purchased. The same license key file should be uploaded and activated on each server you have purchased a license for. Each time you activate the license file the license activation count will increased, until the total number of licenses that the key file contains has been activated. If FusionReactor reports that you have *'no Activations remaining'* on this Multi-Key file and you believe this to be an error, please contact our support team.

Also See

FusionReactor Purchase Page

License Policy

Each physical box (server) requires one FusionReactor license - all of the instances which are installed on that server are covered by the one FusionReactor license. Note, that on a single box you may have multiple instances (e.g. of ColdFusion, Tomcat, JBoss) installed on it, but you still only need 1 license for that physical box. If you use Virtual Machines (VM's) then these are also classed as separate physical servers and therefore require 1 license per VM. For a breakdown of the latest prices, please refer to the buy page.

See Also

FusionReactor Purchase Page

Quick Start to Troubleshooting Server Problems with FusionReactor

This document helps you configure FusionReactor. There are many different scenarios in which FusionReactor is used, each with its own conditions, and each requiring a different configuration. What we aim to do in this document is cover sample configuration for short term (immediate) troubleshooting and longer-term production monitoring.

It's not possible to come up with a universally-perfect configuration, because every scenario is different. What we've tried to do in this document is cover the most important points to consider when configuring FusionReactor for different scenarios.

Getting FusionReactor

Configuration and Diagnosis in Emergencies

Configuration and Diagnosis for Production

Configuration for Enterprise and High-Availability Environments

- Setting Up A Separate Monitoring Instance
- FusionReactor in High Availability Environments (HAE)
- Enterprise Dashboard Servers and Groups

Using log files to troubleshoot issues

Getting Help

Logs

- Stack Trace Log (stacktrace.log)

Getting FusionReactor

We are assuming you've installed FusionReactor already; if not, go ahead and install it now. The installation process is easy – in almost every case the installer package will take care of all the details. We've prepared a comprehensive **Installation Guide** which covers all the details:

http://www.fusion-reactor.com/fr/downloads.cfm

This section helps you configure FusionReactor to diagnose a crashing server.

Configuration and Diagnosis in Emergencies

> **NOTE**
>
> When a J2EE (ColdFusion, JBoss etc) crash occurs, **before restarting, create a log set**. If, after restarting, FusionReactor rotates the log files, data may be lost. The files in the FusionReactor log folder should be zipped up, or at least copied and saved. Your J2EE log files (typically ColdFusion/logs/*) and container log files (/ColdFusion8/runtime/logs) should also be saved.

Configurations

FusionReactor has several settings pages that allow you to do things from selecting how much metric data to view on the Y axis to defining what percentage of memory usage you consider critical. Below is a list of configuration guides for all these types of settings that can help you diagnose problems.

- Request Settings
- Enterprise Settings
- Metrics Settings
- Resource Settings
- Crash Protection Restrictions
- Crash Protection Settings
- JDBC Settings
- JDBC Stack Trace Filter
- Compression Settings
- MIME Type Restrictions
- Exclude URLs
- Log Settings
- Search and Replace
- Filter Restrictions
- Filter Settings
- Settings
- Filter Restrictions

Instant Diagnosis

FusionReactor stores a number of logs in memory, in order to produce useful metrics and graphs. These metrics will be your first point of observation to address issues on an ailing server. The **Resource** and **Request** logs are important to diagnose server problems, since they record the memory/CPU usage of the system and the running requests respectively. They are potentially very useful in an unstable environment, where restarts will cause FusionReactor to lose its in-memory data. All logging within FusionReactor is a computationally inexpensive operation, the limiting factor being available disk space. For detailed information about how the logging system works, please read the Overview of FusionReactor Logs.

Memory Observation

Memory should be your first observation. If the **blue** chart of the **Memory Usage** bar is **consistently** near the top of the graph, consider making more memory available to the J2EE (ColdFusion) process. This can be altered within the **ColdFusion Administrator** in the section **Server Settings -> Java and JVM -> Maximum JVM Heap Size**. If there is insufficient memory on the system to increase the JVM memory, consider reducing the size of the **Template Cache** and **Cached Queries**.

A **sawtooth pattern** in the blue memory graph is normal. This shows Java periodically garbage collecting objects. When **used** memory begins to approach the **allocated** memory value, you may see one or more sawtooth (garbage collection) patterns, as Java attempts to reclaim memory before asking for more. If insufficient memory is reclaimed, you will see the white **allocated** memory bound increase, as Java demands memory from the operating system. If the blue portion steadily rises over the course of an hour, this often indicates a **memory leak**. These are becoming more common as the complexity of J2EE applications increases.

If the used memory (blue) graph is growing rapidly, you can try to ask Java to perform garbage collection yourself by clicking on the **trash can** icon on the Memory Usage graph. However, Java's memory algorithm is very sophisticated and it's unlikely a manual collection will have any significant effect.

CPU Observation

CPU Usage is another useful metric. If the instance is consistently busy (see the blue line on the **Memory Usage** chart) with low load (see Request Activity graph), this might indicate a problem with the pages being run (see Request History), such as infinite loops or runaway queries (consider using the FusionReactor JDBC Driver Wrapper to analyze and prevent JDBC problems).

The Metrics -> Longest Requests page shows the longest running requests. FusionReactor also flags long-running requests with an appropriate label in the **Request History** table. You can configure what FusionReactor considers a long-running request, and how many to store in the Longest Requests table, in the M

etrics Settings page.

Configuration and Diagnosis for Production

When we say Production we mean a stable system which is serving live requests. These systems aren't used to test new applications, and they usually settle into a steady state over time. Once they stabilize, we can perform some simple analysis to tune FusionReactor. The principal of tuning a production system is to watch for anomalies in the performance of the system, in essence defining exactly what constitutes an “exception case”. Once we've identified this, Crash Protection rules can be defined to report when this case occurs. There are essentially two methods of defining these cases: **gut feeling and numerical analysis**. We'll give an example of both, and how they apply to each of the Crash Protection settings.

Gut Feeling

The “gut feeling” principal is best supported by the System Metrics graphs. With the system running steadily for at least an hour, the graphs – when switched to **Hour** mode - clearly show how the system is being used.

Timeout Protection and Request Protection

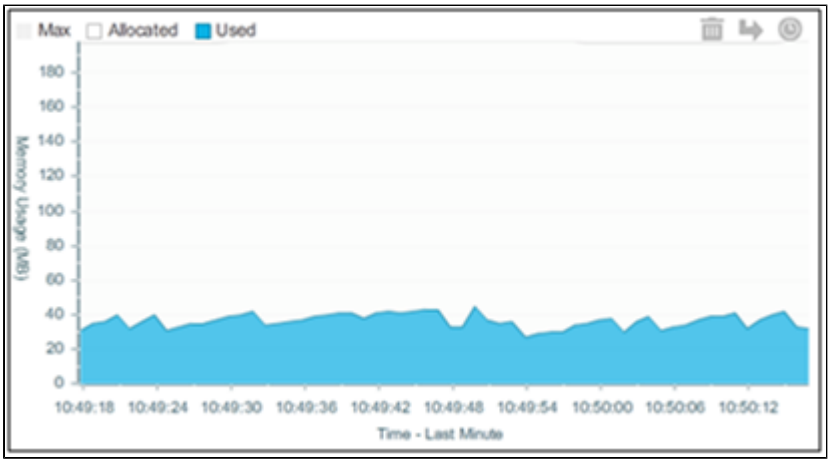
You can use the system metrics graph to check the following:

- Timeout Protection, using the Average Request Time Graph.
- Request Protection, using the Request Activity Graph.

Memory Protection Example

As an example, we set up a load test to simulate a fairly constant load on the system.

You can see in the System Metrics page that the memory displayed in the Memory Usage graph is reasonably constant. If we observe the system for an hour, we would get an even better picture of how the memory is used. For even longer periods, we can set the Resource Settings -> Resource Sampling interval to 60000ms (1 minute), with a long history size. We can then view this data in the Resources -> Memory Graphs -> Memory Heap page. Using a longer time span allows us to get a picture of time-dependent fluctuations such as peak load and scheduled tasks.



Using the data on the Memory Usage graph, we can see that the steady state for our test is around 40MB. Allowing a 100MB overhead for scheduled tasks and peak-time usage, this would give an acceptable operational peak of 140MB, with a remaining “emergency” margin of 60MB. We might then set Crash Protection Memory Protection to 30% (60MB margin = 30% of the total memory).

The system should then be left in Notify mode for a period of time. While running in this mode, adjustments can be made to the protection value based on the notification email. When you are satisfied that the value you have represents the upper margin for the system, change the notify mode to Abort and Notify or Queue and Notify.

In the detailed example above, we've allowed a margin in the operational memory estimation to account for periodic rises in website usage, and scheduled tasks. Scheduled tasks often run at periods of low demand, and are typically resource-intensive. Search indexing, reporting and cleanup operations often run in this batch mode, and often consume memory and CPU cycles. These tasks also tend to run for long periods of time.

Once you have a picture of these tasks, you can add them as CP Restrictions (Crash Protection -> CP Restrictions). Ensure that CP Settings -> CP Restrictions -> Restrictions is Enabled, and the mode is Ignore Matching Requests. CP will then ignore any requested pages set in the CP Restrictions page.

Once you have added your scheduled tasks as restrictions, the operational values for Timeout, Memory and Request Protection can be adjusted, since they will no longer apply to the exceptional cases.

Numeric Analysis

The second method of setting Crash Protection values uses the log files generated by Crash Protection to generate a measure of the statistical spread of page runtimes. Using this measure, we can compute a statistical upper bound to the runtime, which can serve as a starting point for further tuning.

Computing Standard Deviation of Page Runtime

Computing the statistical standard deviation of page runtimes gives you a weighted measure of the spread of the values. We can then use this measure to set the **Timeout Protection**.
The simplest method of computing this value is to take advantage of FusionReactor's **Request Log**. You can get more information on how to use the logs in Overview of FusionReactor Logs.

Computing Standard Deviation of Page in OpenOffice

1. Ensure **Requests -> Request Settings -> Request Logging** is **Enabled.**
2. Allow the system to run in its steady state for a while (the longer the better).
3. Load the **request log** (/FusionReactor/instances/NAME/log/request-0.log) into OpenOffice Calc
    a. Open the file as a **Text CSV** file, with options: **Separated by: Space**, **Text delimiter: "**
4. Create a new sheet (**Insert -> Sheet**), naming the new sheet **Sheet2.**
5. Switch back to **Sheet1**, where the data is, and select **Edit -> Select All**.
6. Select **Data -> Filter -> Standard Filter**, setting **Column F** equal to **COMPLETED**. This will remove any other request statuses from the list. Click the **More** button*,* and select **Copy Results To**. Enter **$Sheet 2.$A$1** in the destination box and click **OK.**
7. When the copy is complete, you should have only the **COMPLETED** entries in Sheet 2.
8. In a free cell which is not in column , enter the formula: **=AVERAGE(L1:L65535)+STDEV(L1:L65535)**
    a. This value is the upper limit at **one standard deviation away from the average.**
9. Statistically,this may still produce too many false positives for Crash Protection, so we may want to use 2, 3 or even 4 standard deviations:
    a. **=AVERAGE(L1:L65535)+4*STDEV(L1:L65535)**

As with the gut feeling approach, this number can only ever be a starting point for tuning: you are bound to get some exceptions, and these should help you decide how to adjust this value.

Configuration for Enterprise and High-Availability Environments

FusionReactor Enterprise Edition makes monitoring more than one server instance easy with the **Enterprise Dashboard**. Any FusionReactor Enterprise Edition installation is able to report the state of the server remotely to any other FusionReactor Enterprise Edition server.

We are often asked about the best configuration for enterprise environments: this section aims to provide some guidance about how best to set this up.

Setting Up A Separate Monitoring Instance

FusionReactor in High Availability Environments (HAE)

Setting Up A Separate Monitoring Instance

By default, FusionReactor Enterprise Edition is configured to monitor itself. You'll see the local instance in the **E nterprise Dashboard**. We don't recommend running any scripts or using the **Server Shutdown/StartUp Alerts** functionality in this mode, since we can't guarantee the behavior of a system monitoring itself.

In Enterprise or High-Availability Environments, we recommend using the FusionReactor Administration Manager (FRAM) as a separate monitoring solution. When you install FusionReactor you will automatically get a FRAM instance installed for you but in Enterprise or High-Availability Environments you should create a separate J2EE container containing **only** the FusionReactor Administration Manager, whose sole purpose is to monitor other instances. This solution isolates FRAM from problems occurring in other J2EE (ColdFusion) instances. Because FusionReactor is licensed on a per-machine basis, no extra licenses are required in this scenario.

We recommend installing a new Tomcat instance for this solution. Simply add the monitored machines to the **M anage Servers** page and use the **Enterprise Dashboard** to monitor the status. If more detailed information is required, click on the small **arrow** icons on the dials or the server cube (see figure 1): you'll be logged into the remote server automatically and directed to the appropriate detail page.

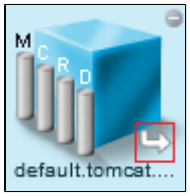


Figure 1

FusionReactor in High Availability Environments (HAE)

In a High Availability Environment (HAE), we recommend a redundant Dashboard configuration: FusionReactor is installed on two (or more) physical servers. Each monitor is configured to monitor all the cluster instances (or a subset of them) **and the other monitor**. Using **Server Shutdown/StartUp Alerts**, the system administrator can monitor the status of all the affected servers. For maximum protection, each FusionReactor monitor should be installed in its own Tomcat server.

If you configure FusionReactor in this way, you should be careful how you handle **Enterprise Scripting** if you are using it to perform restart functions.

1. Never use scripting to restart the monitoring instance. Because of the order that FusionReactor's subsystems start and stop, a normal restart may cause a script to fire.
2. Always offline servers in the Dashboard before performing scheduled maintenance. Otherwise, Enterprise Scripting may try to restart them.
3. Do read the document FusionReactor Enterprise Scripting, which is available on our website and is installed along with FusionReactor.

Enterprise Dashboard Servers and Groups

Another recommendation is for people with large numbers of servers. The **Enterprise -> Manage Groups** page allows you to group together monitored instances into logical groups. You can use this functionality to partition the Dashboard into logical areas that fit your organisation. There's no limit to the number of servers that can go into a group.

If you need an instance to be present in more than one group, simply add it again using **Manage Servers**, then put it in the second group. Provided the **URL** used to access the monitored instance is the same, FusionReactor will only poll the system once, and make that data available in both Dashboard icons.

Once your servers have been organised into groups, you can then simply monitor the **group cube** – it will

always reflect the **worst status** of the group. See for figure 1 for an example of this.

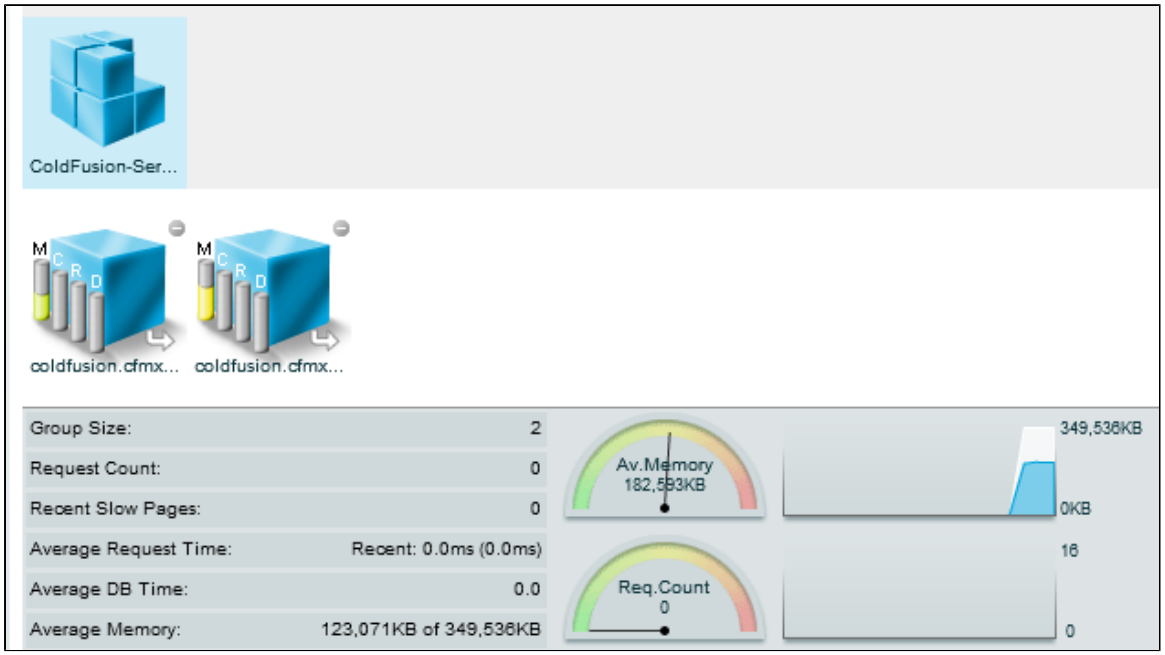


Figure 1

Using log files to troubleshoot issues

FusionReactor provides many helpful graphs and metrics to help you see how your system is performing. These graphs are available in real-time and will also show you a short history of what happened (depending on the amount of data which has been generated), but what happens when your server begins to crash? Lots of this data is stored in-memory so it can be used to generate reports and lists for the FusionReactor Administrator. If you are having problems which cause (or necessitate) restarts of the server process, this in-memory data is lost - however all of this data will have been captured to one of FusionReactor's extensive log files.

This data can in almost all cases be used in a post-mortem session to pinpoint the location of problems within a system. In addition to immediate post-mortem work, the logs can also be used to analyze usage patterns, trends and to derive longer-term aggregate data.

Also See

Importing and Graphing Data in Excel
Methodology in OpenOffice, where we used the **request log** with standard deviation to derive a nominal page run time value.

Getting Help

Fusion Products Knowledge Base

We provide a knowledge base full of technical articles and helpful hints at the following URL:

http://www.fusion-reactor.com/fr/support.cfm

From here you can search the support database, see what's new, what our hot issues are, as well as finding online versions of our printed manuals.

Google Group

We also provide a Google Group for the community to ask questions in.

http://groups.google.com/group/fusionreactor

We check and reply to this group regularly, so it's also worth searching to see if your question has already been answered.

Support Mailbox

You can also mail us FusionReactor questions at:

support@fusion-reactor.com

We can also provide software consulting services if you're having problems with your application:

http://www.fusion-reactor.com/support/supportservices.cfm

Logs

Stack Trace Log (stacktrace.log)

About this log

Every entry that is written to this log file contains a Stack Trace of all running Threads on the JVM that your server is running on. Each entry provides a very detailed snapshot of the state of your server at that specific time, this information can be used to aid you diagnose issues that you are having with your server.

The log file is formatted as follows:

Field	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date/Time (ms)	3 [C]	This is the millisecond time at which this entry was written
Requester	4 [D]	Whatever requested the log entry e.g. Crash Protection
Stack Trace	5 [E]	This is a Stack Trace of all running Threads on the JVM (that your server is running on) at the time this entry was written. This field consumes a lot of space as the information here is typically used for diagnostic purposes.

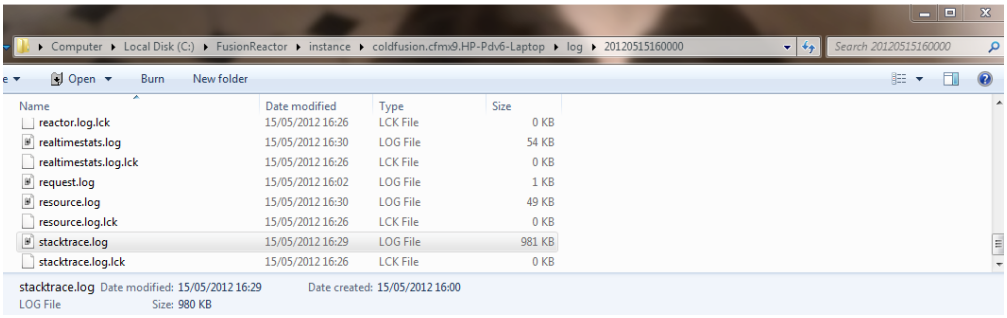
When are entries added to this log file?

- When Crash Protection is triggered.
- When you navigate to Resources > List All Threads and then click on the button that contains the text "Stack Trace ALL"
- When you navigate to Requests > Running Requests and click on the button that contains the text "Stack Trace ALL"

- When you navigate to Requests > Running Requests, navigate to a running request and click on the this button,  .

Where can I find my "stacktrace.log" file?

Navigate to "<FusionReactor_Home>/instance" where "<FusionReactor_Home>" is the root directory of your FusionReactor installation. The instance directory contains a list of directories named after your FusionReactor instances. Navigate to the directory that is named after the instance that you are interested in and then navigate to the "log" directory. The log directory contains a directory that is a series of numbers, this is the epoch time that this directory will be subject to log rotation. Navigate to this directory, this is where the stacktrace.log file can be found.



Crash Protection and Monitoring with FusionReactor

FusionReactor 4 encompasses several technologies designed to ease the burden of J2EE (including ColdFusion) support organizations – two of which are **Enterprise Monitoring** and **Crash Protection.**

Crash Protection is, of course, one of the important features of FusionReactor, and has been present since the initial release. FR4 builds on this technology and brings advances in **Restrictions** (which requests will be subject to Crash Protection) and we have also extensively reviewed the mechanisms FusionReactor uses to manage and control requests. Server downtime is often measured in dollars, not minutes, and Crash Protection continues to be a focus of the FusionReactor team.

Enterprise Monitoring (available to customers with enterprise licenses) gives you the ability to visualize the state of a machine cluster, together with behind-the-scenes uptime monitoring, and the ability to send mail when a machine becomes unavailable. Easy-to-read instrumentation including color-coding, realtime graphs and bars, together with advanced grouping, server-offlining and user-definable warning and critical thresholds make our **E nterprise Dashboard** an essential component for hosting providers, or customers with a cluster of J2EE (e.g. ColdFusion) installations.

The Enterprise Dashboard is fully explained in its own section, so we won't be covering it here, except where necessary to illustrate how a particular issue is visualized.

Scenarios and Scope

This documentation will cover the usage of FusionReactor to ensure the availability of your systems. We'll cover Crash Protection and Enterprise Monitoring in some detail. We'll show you how to use FusionReactor to cross-monitor a cluster of machines, such as is typically found within companies with fail over solutions, or in hosting providers running large farms of machines.

We'll illustrate in detail, using worked examples, how FusionReactor's sophisticated Crash Protection works to keep your system up and alive, even in the face of an onslaught of requests, bad queries, pages that take too long and memory which is approaching upper margins.

Although every installation of FusionReactor is different – and we're constantly amazed at the diverse and interesting environments FusionReactor is deployed within – we'll also show you how you might tune Crash Protection to optimize your production environments.

Finally, we hope you are happy with FusionReactor. We're convinced it will provide a significant return on your investment, maybe in some cases even paying for itself by preventing one outage. Over multiple version releases we have made significant advances in enterprise monitoring, protection and visualization, but a lot of it would not have been possible without feedback from our customers. We're always grateful for comments on our products; if you have any criticisms or ideas, please send them to: support@fusion-reactor.com

Crash Protection Basics

- Design Goals
  - Common J2EE Problems
  - Types of FusionReactor Crash Protection
  - Survival Strategy
- Crash Protection Implementation
  - Implementation
  - Abort Strategies Explained
  - Risks and Restrictions of Request Abort
- The Crash Protection Restrictions Engine
  - Crash Protection Restrictions Engine
  - Customer Demographics
  - Restrictions Engine Rules

Crash Protection In Action

- Test Pages
- Exercising Timeout Protection
- Exercising Memory Protection
- Exercising Request Quantity Protection
  - Using the 'Queue' Strategy
- Using the Redirect Method
  - Best Practices in Redirection

Understanding Crash Protection Restrictions

- Modes
- Rule Basics
- Examples of Restrictions
  - Excluding Batch Jobs
  - Including Specific Hosts
  - Including a Specific Action Page

Crash Protection Alerts

- Email Alerts

Crash Protection Basics

Design Goals

- Common J2EE Problems
- Types of FusionReactor Crash Protection
- Survival Strategy

Crash Protection Implementation

- Implementation
- Abort Strategies Explained

- Risks and Restrictions of Request Abort

The Crash Protection Restrictions Engine

- Crash Protection Restrictions Engine
- Customer Demographics
- Restrictions Engine Rules

Design Goals

Common J2EE Problems

Types of FusionReactor Crash Protection

Survival Strategy

Common J2EE Problems

Crash Protection attempts to keep your server alive by watching for and preventing (or minimizing) the effects of three scenarios, which we identified as being common causes for outages:

- Requests taking too long are often an indicator of poor code design. In interactive systems like the web, requests must run in as short a time as possible. It's often preferable to receive an error message rather than continue waiting for a request which may never complete.
- Too many requests running simultaneously often leads to resource starvation, and although it's usually possible to tune this using the J2EE engine itself, the system must usually be restarted for it to take effect – and the option is usually well hidden. The majority of J2EE engines – including Adobe's JRun – control resource access using thread pools, and having too many requests running at once can reduce or exhaust these pools to the state where it's no longer possible to recover the system, leading to a costly restart.
- Exhausted memory occurs very frequently in production environments. On shared installations as well as dedicated servers there never seems to be enough memory. Approaching memory margins is a sign that the server is overloaded, or one or more requests are creating a lot of objects and using more than their fair share of storage.

Types of FusionReactor Crash Protection

We designed three crash protection rules to counter these common situations. There's no need to alter any J2EE engine settings to take advantage of these rules, and no container restart is required for them to take effect.

- **Timeout Protection:** Any request which runs longer than a given threshold causes this rule to fire.
- **Request Quantity Protection:** once the number of simultaneous requests reaches this threshold, further incoming requests cause this rule to fire.
- **Memory Protection:** once memory has breached this threshold (specified as a percentage of total), further requests cause this rule to fire.

Using these rules, either individually or in combination, FusionReactor can keep a server alive in a marginal situation for much longer than it would otherwise be available. In the vast majority of cases, Crash Protection rules can keep a server up and responding long enough for the marginal situation to clear, allowing processing to continue normally.

Survival Strategy

Along with each rule, it's possible to specify how FusionReactor will react when that rule fires - the Survival Strategy.

The three available survival strategies are:

- **Abort and notify** – once the rule fires, mail is sent to the administrator, and further incoming requests are aborted either with a fixed message, or a URL redirect to a custom page.
- **Notify** – once the rule fires, mail is sent to the administrator, and the request is allowed to continue normally.
- **Queue and notify** – once the rule fires, mail is sent to the administrator, and the request enters a holding queue until the marginal situation is resolved. The length of the queue – by default 60 seconds – can be configured.

To avoid a flood of email in marginal situations, notification can be turned off completely, or can be set to only send email once in a given period – by default one minute. All three survival strategies are available for each of the three rules, except **Timeout Protection**, which can't use the **Queue** rule, since the requests it monitors can't be queued once they've started.

Crash Protection Implementation

Implementation

Abort Strategies Explained

Risks and Restrictions of Request Abort

Implementation

FusionReactor is implemented as a J2EE Filter and if installed correctly, should be the very first filter in the J2EE chain – this is how the install package installs it, and how our manual installation guide explains it too.

Being the first filter in the chain allows FusionReactor to control exactly what runs within a given J2EE server, and provides the maximum control during marginal situations.

FusionReactor uses its own web server and threading pools to ensure that it does not use J2EE container resources, and therefore also does not rely on an external web server to be available – something that is often one of the first things to become unavailable when resources get tight.

We explicitly designed all aspects of FusionReactor to consume as little time as possible. Turning FusionReactor on should have very little impact on your J2EE application. Part of the product QA cycle involves a multi-day run test under very strenuous conditions to check there are no performance or memory problems.

Activating Crash Protection rules, therefore, can be a routine action; even if you only require a notification and no explicit survival action.

The **Memory Protection** rule is evaluated before the **Request Quantity** rule. If a situation exists in which both rules should logically fire, the **Memory Protection** rule (often the more dangerous of marginal situations) is given preference to deal with the situation.

Abort Strategies Explained

Apart from the **Notify** action, which sends a detailed email during Crash Protection interventions, there are two other major strategies for dealing with requests: **Abort** and **Queue.**

Queue

This strategy attempts to alleviate the marginal situation by temporarily pausing incoming requests until the condition which caused the rule to fire no longer exists.

The length of the queue can be specified, after which requests will be terminated regardless of the survival strategy.

There is no limit on the size of the queue, so if a large quantity of requests are present they will consume tracking resources inside FusionReactor (albeit temporarily). This strategy is therefore best used on systems

with sufficient memory, or on systems where the volume of requests is known and not expected to become prohibitive.

Abort

This strategy aborts requests. When used with the **Memory** or **Request Quantity** rules, it is best imagined as a 'reject', rather than an abort. The request is summarily rejected and not allowed to proceed inside the J2EE engine. The abort strategy – redirect to URL or display of fixed message – is applied prior to the abort action.

When used with the **Timeout Protection**, FusionReactor uses strong thread manipulation techniques to make sure requests are stopped.

In all abort actions, FusionReactor will process the terminating requests for statistical and tracking purposes.

Risks and Restrictions of Request Abort

JVM Thread Aborts

Because a request abort is a last-ditch effort to prevent a marginal situation become critical, FusionReactor attempts to use a rather strong method to halt processing.

In the vast majority of these cases, the thread stops immediately and FusionReactor is able to recover the system.

However, in a very small set of theoretical cases – and we must emphasize we've never seen this happen during innumerable test runs – the JVM could be restarted. Most J2EE service wrappers will automatically restart it in this case, causing a minor outage – which would in any case be less serious than a hung system which would otherwise occur.

In the light of our test results, and the use of request abort in our own production environments, we are still utterly convinced that the benefits of this strategy far outweigh the infinite simally small risk.

Native-Bound Threads

Due to restrictions in the JVM, it is not always possible to immediately stop a thread. If a thread is currently engaged in a blocking native operation, i.e. performing I/O in a JNI method (sockets are a good example of this) then the JVM will not be able to abort the thread until it completes.

For this reason, you may not see requests disappear immediately when aborted with Crash Protection or killed manually from the FusionReactor Administrator.

We're continuing to investigate this problem – which is a restriction of the Java virtual machine – and will release an updated version of FusionReactor when we have a solution.

J2EE Containers and Dead Threads

After aborting a thread, FusionReactor processes the requests for statistical and display purposes inside the FusionReactor Administrator, then disposes of the original thread object by returning it to the J2EE engine.

Although the object no longer represents a valid thread, some engines do not check this,assuming the thread to still be runnable, and return the object back to their internal thread pool.

In some cases, this object can then be picked by the pool to run a new incoming request. This will then fail immediately, at which point the engine will remove the thread from the pool. The client then sees an error message, usually accompanied by a HTTP status 500 – Internal Server Error.

Again, this happens infrequently and is vastly preferable to a total outage.

The Crash Protection Restrictions Engine

Crash Protection Restrictions Engine

Customer Demographics

Restrictions Engine Rules

Crash Protection Restrictions Engine

This section illustrates a brand new feature to FR3 – **Crash Protection Restrictions.**

This powerful filter is used to configure exactly which requests will benefit from Crash Protection. There are certain classes of request – batch jobs, database update, data warehousing – which are known to consume time and resources and which could be safely ignored by Crash Protection, so-called 'false positives'.

Customer Demographics

During the design of this feature, we identified two key demographics for whom this would be an important technology:

- The small homepage or small-company site which is agile and which will protect all requests **except** its long-running updates
- The large enterprise site, possibly with many web services and non-interactive services, which will be default **not** protect requests, preferring to define small groups of interactive pages which will be protected.

Instead of attempting to aggregate these groups into those who would and those who wouldn't use the feature, we made the Restrictions Engine reversible. The engine – which is totally optional – can be run in two distinct modes:

- All requests are protected **except** those matched by a filter rule (exclude mode).
- All requests are ignored **except** those matched by a filter rule (protect mode).

We feel these two modes accurately reflect the needs of the two demographics above.

Restrictions Engine Rules

In order to decide whether a given request will (or will not) be protected, the Restrictions Engine evaluates a number of user-defined rules.

Each rule allows fine-grained specification of exactly which requests it matches.

A rule is specified as an exact match or a Regular Expression as defined by the relevant Java JRE specification1. As well as the URL path, the following optional components can be used to define the rule:

- The **requested hostname** – useful to differentiate requests on multi-homed systems
- The URL parameters – useful for requests whose behavior changes according to their parameters

Additionally, each rule can specify whether pages matched by that rule will be protected by (or excluded from) all protections or just Timeout protection. If the list is running in exclude mode, in which all requests are protected by default, it's also possible to specify whether any excluded requests will still be tracked for statistical purposes (e.g. runtime tracking etc.).

Crash Protection In Action

Test Pages

Exercising Timeout Protection

Exercising Memory Protection

Exercising Request Quantity Protection

- Using the 'Queue' Strategy

Using the Redirect Method

- Best Practices in Redirection

Test Pages

In order to exercise the three types of Crash Protection, we have written four pages (source code is presented in the appendix):

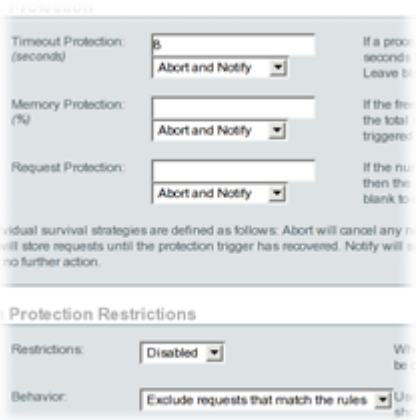
- **pause.cfm** – takes a 'timeout' parameter (in seconds). This page simply sleeps for the given number of seconds before completing. It's used to exercise the **Timeout Protection**.
- **load.cfm** – takes a 'load' and 'timeout' parameter. This page loads the system by spawning itself to produce the required number of requests. Each page then sleeps for the given number of seconds before completing. It's used to exercise the **Request Quantity Protection**.
  E.g. Calling *load.cfm?load=5&timeout=10* would cause 4 requests to be spawned (giving a total load of 5), and each one would sleep for 10 seconds.
- **grow.cfm** – takes a 'freemem' parameter (MB) and allocates memory continually until that threshold is reached, then sleep for 60 seconds.
  E.g. Calling grow.cfm?freemem=10 would instruct the page to allocate memory until there was only 10 MB remaining, then sleep for 60 seconds.
- **system.cfm** – simply outputs the version of Java in use. This page is used as a 'hello world'-type request to see if the system will accept requests.

These pages are placed in a temporary folder in our test machine. We use the following URL to access them:

http://int0234/tmp/ts1/**pause.cfm?timeout=10**

Exercising Timeout Protection

In the FusionReactor Administrator, access the CP Settings area by selecting it in the contents area. You may need to open the 'Crash Protection' section to see this.



Once the crash protection settings page has loaded, ensure the Restriction engine is disabled, and timeout protection is set to 8 seconds, with the **abort and notify** strategy. Illustration 1 shows this setting.

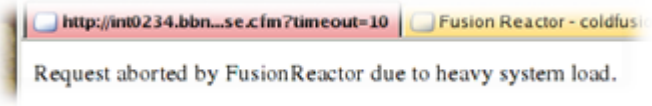
Once this has been committed with the 

button, the Timeout Protection becomes active. We can then run the pause script to show the protection in action.

In our browser, we use the url **pause.cfm?timeout=10** to ensure that the script is still running at the timeout protection limit of 8 seconds. In fact, since this is the first time the page has been run, ColdFusion will require a short period to compile it first. This additional time is also counted and tracked as part of the 8-second

protection limit. The page would therefore run for longer than the 10-second pause time, giving the Timeout Protection ample opportunity to activate.

When the protection activates, the **Display Message** abort strategy is used, and the user sees an appropriate message.



The abort strategy and the display message are both configurable using the Crash Protection Settings screen, which we used earlier to set up the protection. Our **pause.cfm** script simply uses a Java Thread.sleep() call in order to go to sleep, so it is easy for FusionReactor to wake up this thread prior to killing it. As we noted earlier though, if the request was performing socket I/O, FusionReactor would be forced to wait until the script returned from the socket operation before it could terminate it. Most operations can be terminated immediately though. The **Request History** page shows the abort by annotating it with a 'timeout' flag, as illustrated X.



Exercising Memory Protection

Exercising memory protection is a little more difficult, due to the disparate memory configurations in use, but here's how we use FusionReactor to show it in action in our test environment.

The first step is to work out how much memory is available to us. Instead of hunting through J2EE configurations, we can use FusionReactor to tell us.

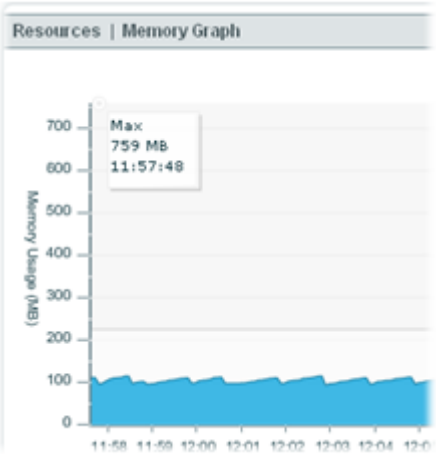


Figure 1

Access the **Memory Graph** (accessible from the table of contents under the 'Resources' section) and have a look at the figure at the top left of the Y axis (see Figure 1 for our value – about 759 MB).

Memory Protection uses a percentage value as the upper limit, to ensure a marginal situation does not become critical. If we define a Memory Protection of 10%, this translates to about 76 MB. The actual effect of this protection will be to reject new requests when the quantity of free memory drops below 76 MB. During this test, we'll be using the **Memory Graph** to observe the situation.

In the Crash Protection Settings page, we set up the **Memory Protection** to activate at 10%, and make sure any value in the **Timeout Protection** box is removed. We commit the change using the



button, as before.

The system is now primed for the test. We use the **grow.cfm?freemem=8** script to run the system to 700 MB total memory demand (about 60 MB free memory), a value which is well within the 76 MB margin (which occurs

at 686 MB).



Figure 2

Any requests which occur while the system is in this marginal situation will be rejected by the **Memory Protection**. We can watch the **Memory Graph** to observe the memory load during the test. When the **grow.cfm** page reports it has exhausted memory to the required limit, we run the **system.cfm** script, which is immediately rejected (See Figure 2). Figure 3, below, shows the memory graph during the test. Shows the memory load during the test, including the calm period before the test, and the demand peak caused by **grow.cfm**.

The sawtooth pattern before and after the test is a classic example of Java's periodic garbage collection in action. The memory demand caused by **grow.cfm** isn't actually released back to the system until a few seconds after the request completes – again, this is normal since Java only reclaims this memory when it's urgently required, or in this case, when system activity declines to a calm level. By observing the **Request History** page (Illustration 6, below), we can see that FusionReactor correctly rejected the **system.cfm** request during the marginal situation, and that it's been tagged with a 'rejected' flag.

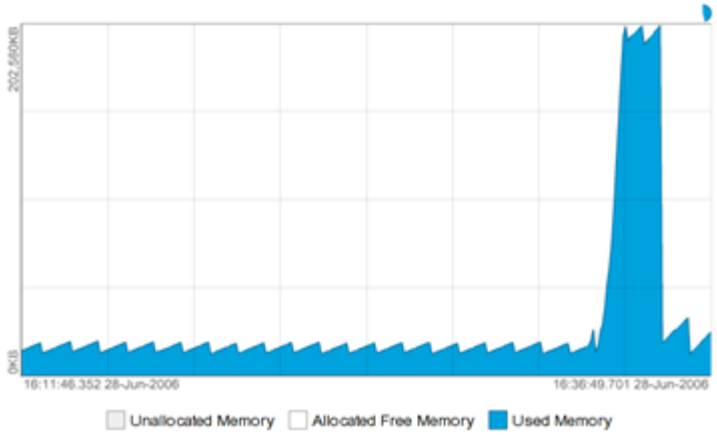


Figure 3

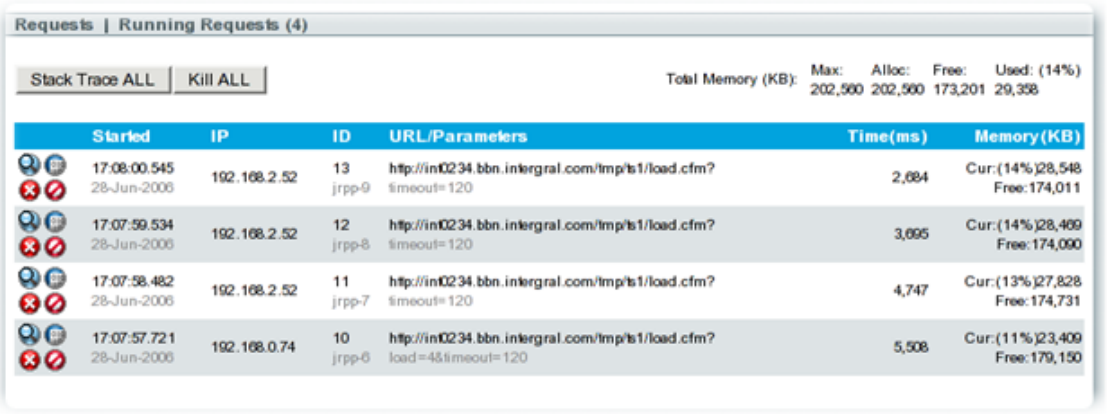
Exercising Request Quantity Protection

The Request Quantity Protection rejects incoming requests when the total request load on the system exceeds a certain value. This protection is completely independent of any values defined by your web server or J2EE container. Like other FusionReactor protections, this value can be changed without causing a restart of any other software components, minimizing outage time.

To exercise this protection, access the Crash Protection Settings page using the table of contents, and remove any other protections which may be left over from previous tests.

Set a **Request Protection** value of **4** and commit the changes using the 

button. When the number of running requests rises to 4, FusionReactor will control the load by rejecting any incoming requests.



Using the **load.cfm?load=4&timeout=120** script, we load the system with 4 concurrent requests for 120 seconds, visible in the **Running Requests** page, above.

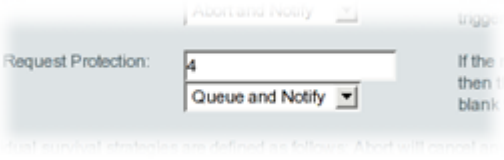
Running the **system.cfm** page while this load exists causes the page to be rejected.



Using the 'Queue' Strategy

We will also use this scenario to demonstrate the 'queue' survival strategy. Any requests which arrive are queued by FusionReactor to give the system time to recover. The default queue period is 60 seconds, which is usually ample time.

Use the **Crash Protection Settings** page to change the survival strategy for **Request Protection** to **Queue and Notify.** We'll leave the queue timeout at 60.



Firstly, we'll show requests being queued and then completing. Running the **load.cfm?load=4&timeout=60** scri pt, we increase the system request load to the Crash Protection limit of 4. Running the **system.cfm** we see the page pause until the **load** pages complete and the total load reduces below the margin.

If a page remains in the queue longer than the **Queue Timeout** period, it will be rejected by the protection. We can illustrate this by running the **load.cfm?load=4&timeout=120**, and then running the **system.cfm** script. The load will be applied for 120 seconds, which is longer than the queue period. The **system.cfm** script will then be **expired** from the queue. The page expired from the queue after 60,797ms.



Using the Redirect Method

Up to this point, we have been using the abort strategy to return a short message to the client whenever a marginal situation arises. While this is is usually sufficient to warn a client that the system is currently recovering, there are cases where customers may want to return their own branded pages to the user.

The Crash Protection system therefore also allows an HTTP **Temporary Redirect** message to be sent to the client. This is configured on the **Crash Protection Settings** page in the **Survival Strategy** section, as shown in Figure 1.

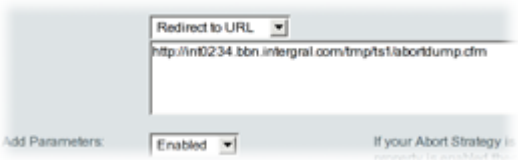


Figure 1

When a marginal situation arises which would otherwise cause a page to be aborted or rejected, FusionReactor will now attempt to send an HTTP redirect to the client. If the 'Add Parameters' setting is enabled, FusionReactor will also supply some information as to the cause of the redirect. This information can be used in customers' own scripts to provide custom branded error messages.

If enabled, parameters are passed as URL (GET) variables and are named as follows:

- DETECTION_METHOD
  One of **timeout**, **freemem** or **requests**, specifying the protection which fired.
- THRESHOLD_VALUE
  The threshold value associated with this protection.
- ACTUAL_VALUE
  The actual value which caused this rule to fire.

For the **timeout** method, the threshold and actual values are specified in milliseconds. For the **freemem** method , the values are in megabytes.

Best Practices in Redirection

We believe redirection is a great way to provided branded error output to clients, but there are a couple of caveats to be aware of.

Making Marginal Situations Worse

If you target a redirection to a page which resides within the J2EE container that FusionReactor is currently monitoring, you may make a marginal situation worse by providing additional load.

Using a redirection target in the same container is possible - indeed this may be your only option - so try to design target pages which place low demands on the server:

- Few images
- Low memory demand from J2EE scripting languages (JSP, CF etc.)
- Low processor demand – no database work, no CPU-intensive computation

Design of redirection targets should therefore be a compromise between corporate branding guidelines and server load.

Redirection Isn't Always Possible

FusionReactor uses an HTTP Temporary Redirect to instruct the client to display the redirection target page. This is the safest, fastest way of sending a redirect. It guarantees the returned page is valid, and produces low server overhead in memory and CPU cycles.

However, this redirection method can *only* be used if the HTTP response headers have not yet been committed. If a page has already begun streaming body data, the thread will be aborted as normal, but the redirection will not occur.

If page design takes this into account, by performing business logic before display logic, this is not likely to become a problem.

Stuck Threads

If your redirection target lies within the same J2EE container which caused the redirection, the problem outlined

in J2EE Containers and Dead Threads (see page 9) can occur. If the J2EE container picks the newly-stopped thread to serve your redirection target without first checking that the thread is runnable, a 500 Server Error can occur.

Crash Protection Restrictions

Finally, you may like to add your redirection target as a rule in the **Crash Protection Restrictions** engine, in order that it will be ignored by Crash Protection if it occurs.

Understanding Crash Protection Restrictions

Modes

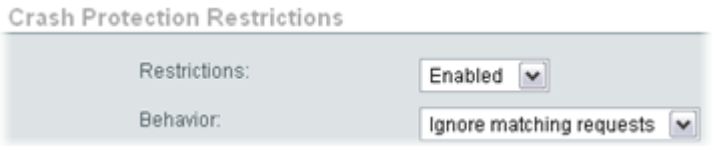
Rule Basics

Examples of Restrictions

- Excluding Batch Jobs
- Including Specific Hosts
- Including a Specific Action Page

Modes

The Restrictions Engine decides which requests will be monitored by Crash Protection, and which requests will be allowed to proceed unobserved. The engine is configured in FusionReactor Administrator in the **Crash Protection Restrictions** section of the **CP Settings** page. Any changes you make to these settings, as well as any changes you make to individual rules, become active immediately without necessitating any software restarts. This allows you to test and tune rules 'on the fly'.



The engine operates in one of three modes:

Mode	Description
Disabled	in which the engine is completely bypassed. In this mode, all requests are monitored by Crash Protection if any protection is currently active.
Ignore matching requests	in which the engine will, by default, protect all requests. Any requests which match a rule are not monitored.
Protect matching requests	in which the engine will, by default, ignore all requests. Any requests which match a rule **will** be monitored.

Engine Overhead

One of the chief design goals of the Restrictions Engine is that it demand very low overhead during the course of a request. The engine is optimized for very low CPU and memory demand and is almost undetectable when active. The Restrictions Engine can therefore be used even on very busy systems.

Rule Basics

Rules for the Restrictions Engine are manipulated using the **CP Restrictions** item in the table of contents. The settings page allows you to then add, remove and change any rules currently entered into the engine.

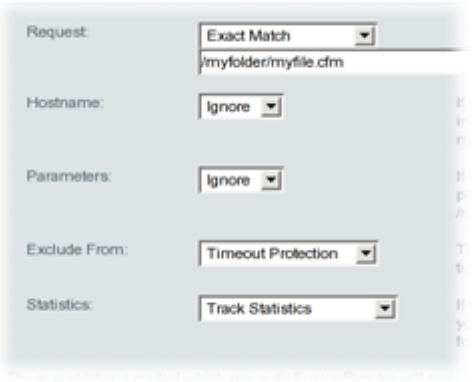


Figure 1

Rules are evaluated from top to bottom, and the first one which fires causes the engine to stop evaluating any further rules.

The Restrictions Engine matches each rule against components of the request URI. The exact components used during the match are selectable using the fields on the rules form (see Figure 1).

The exact options available depend on the current mode of the engine: in **Protect** mode, the **Statistics** option is not available – all requests which match a rule are tracked for timing statistics. The labels used also change to reflect the mode of the system, and the make the meaning of the rule easier to understand.

Request

These two fields define the match mode of the rule, and the actual match string. The drop-down box specifies whether the text field is an exact string match, or whether the field contains a Java Regular Expression.

The remaining options tell the engine exactly which request fields to match against the entered string.

Hostname

If enabled, this field specifies that the string begins with a **hostname**. This allows rules to target specific pages when addressed by multiple sites. The value of this field should match the HTTP 1.1 'Host' header. No scheme (http:// or https://) should be applied. This header is used by web servers, browsers and J2EE containers to differentiate requests for multiple websites which may reside on the same physical system.

A later section will illustrate how to use this feature to protect only certain page requests.

Parameters

This setting specifies whether FusionReactor will differentiate pages based on their URL (GET) parameters. A common design pattern is to change the behavior of a request based on the information provided as URL variables.

An example of this might be a **doAction.jsp** page, whose action is specified as a parameter. Some behaviors of this page should be exempt from protection (**doAction.jsp?action=PopulateDataWarehouse**, for example), while some should be observed and tracked (**doAction.jsp?action=ServeFile**).

This setting allows FusionReactor to treat the two requests separately, and decide whether to protect them based on the URL parameters.

Exclude From

(in **Protect** mode, this field is called **Protection Type**).

Specifies which specific protection is affected by this rule.

If the engine is in Exclude mode, this field specifies which protections any matching requests will be exempt from – either **Timeout Protection** or **All Crash Protection**. If exempt from only Timeout Protection, a matching requests will still be protected by Request Quantity Protection (i.e. will be rejected if the request load is too high) and Memory Protection (i.e. will be rejected if the memory demand is too high). If exempt from **Timeout Protection**, requests will be allowed to run to completion, provided neither the Memory nor Request Quantity protections are activate.

If the engine is in Protect mode, this field specifies which protections will apply to any matching requests – either **Timeout Protection** or **All Crash Protection**. Any requests which don't match will proceed into the J2EE container unprotected.

Statistics

(only available in **Exclude** mode)

If a request matches a rule and is therefore excluded from protection, this setting specifies whether its timing values will still contribute to FusionReactor metrics.

Examples of Restrictions

Excluding Batch Jobs

Including Specific Hosts

Including a Specific Action Page

Excluding Batch Jobs

This example excludes our named batch jobs from timeout protection.

Configuration	Value	
Active Protections	Timeout Protection @ 8 seconds	
Engine Mode	Exclude	
Rule	Exact match	/scripts/CleanUpDatabase.jsp
Exclude from	Timeout Protection	
Page Decisions	/scripts/CleanUpDatabase.jsp?db=MyDatabase	**Ignored**

If we want to exclude this page anywhere it occurs, we can use a regular expression.

Configuration	Value	
Active Protections	Timeout Protection @ 8 seconds	
Engine Mode	Exclude	
Rule	Regular Expression	(.*)CleanUpDatabase\.jsp

Exclude from	Timeout Protection	
Page Decisions	/scripts/CleanUpDatabase.jsp?db=MyDatabase	**Ignored**
	/bigsite/jspscripts/CleanUpDatabase.jsp	**Ignored**

Similarly, if all our batch scripts were named **batch**<job>.jsp*, we could also ignore them with an appropriate regular expression.

Configuration	**Value**	
Active Protections	Timeout Protection @ 8 seconds	
Engine Mode	Exclude	
Rule	Regular Expression	(.)**batch(.)**\.jsp
Exclude from	Timeout Protection	
Page Decisions	/scripts/batchCleanUpDatabase.jsp?db=MyDatabase	**Ignored**

Including Specific Hosts

If we wanted to include a specific host in Crash Protection only, the following rule might suffice (NB the alias “testvm234” points to the same machine as the first URL).
Active Protections
Timeout Protection @ 8 seconds

Configuration	**Value**	
Engine Mode	Protect	
Rule	Regular Expression	Int0234.bbn.intergral.com/(.*)
Protection Type	All Crash Protection	
Hostname	Check	
Page Decisions	http://int0234.bbn.intergral.com/testPage.jsp	**Protected**
	http://testvm234.bbn.intergral.com/testPage.jsp	**Ignored**

Including a Specific Action Page

The following rule specifies an action page with many parameters which normally takes a few minutes to complete when run in a certain mode. We exclude it from Timeout Protection.

Active Protections
Timeout Protection @ 8 seconds

Configuration	Value	
Engine Mode	Exclude	
Rule	Regular Expression	(.)**mightyActionPage\.jsp(.**)action =backupDb(.*)
Protection Type	Timeout Protection	
Parameters	Check	
Page Decisions		**Protected**
		Ignored

Crash Protection Alerts

Email Alerts

How do I receive an email when Crash Protection is fired?

You can configure your FusionReactor instance to send you an email when Crash Protection is fired, the Crash Protection Settings page explains how you can configure Crash Protection i.e. define how you want Crash Protection to work. The Settings page explains how to configure your Email Server so that Crash Protection can send you emails.

What should I expect to see from a Crash Protection email?

Email Title

The email title contains the nature of the Crash Protection, followed by a title to indicate that the email was sent from Crash Protection and finally the FusionReactor instance in which Crash Protection was fired in. Below is a list of all the types of Crash Protection emails that you will ever see.

- [Timeout Protection] FusionReactor Crash Protection Alert [HP-Pdv6-Laptop]
- [Memory Protection] FusionReactor Crash Protection Alert [HP-Pdv6-Laptop]
- [Request Protection] FusionReactor Crash Protection Alert [HP-Pdv6-Laptop]

Email Body

The email body gives you an insight into what was happening in your server when Crash Protection was fired. This information will help you diagnose what happened. In order the information in the email body includes:

- A brief message explaining what triggered Crash Protection and why.
- Information on the request that triggered Crash Protection such as request URL, return status code and client IP address.
- Any action that was taken on the request that triggered Crash Protection, for example you may have configured Crash Protection to abort the request.
- A brief description of all requests that were running on your server when Crash Protection was fired.
- A Stack Trace of all the Threads that were running in the JVM (that your server runs on) when Crash

Protection was fired.

- A detailed description of all requests that were running on your server when Crash Protection was fired such as request URL, return status code and client IP address.

Examples of Crash Protection Emails

Below are some examples of what Crash Protection emails looks like, to make them easier for you to read they have been shortened (mostly the JVM Stack Trace). Simple click the email title to view the email and click the email title again to hide the email.

[Timeout Protection] FusionReactor Crash Protection Alert [HP-Pdv6-Laptop]

```
Hello, this is the FusionReactor Crash
Protection agent at HP-Pdv6-Laptop,
instance: coldfusion.cfmx9.HP-Pdv6-Laptop.
http://HP-Pdv6-Laptop:8088/fusionreactor/f
html.cfm?mode=main
The following alert was generated at Tue
May 15 10:40:00 CEST 2012.
[Notify] Crash Protection - Request Run
Time Alert: Actual value at detection:
1297ms, threshold value: 1000ms

Current (Triggering) Request:
Request ID: 5
 Request URL:
http://127.0.0.1:8500/cf/sleep.cfm
 Status: EXECUTING
 Started (Milliseconds): 1337071199613
 Started (Date/Time): 2012-05-15
10:39:59.613
 Thread ID: web-0
 Client IP Address: 127.0.0.1
 Request Method: GET
 Execution Time (ms): 1301
 Used Memory (percentage): 61
```

Max Memory (KB): 466048
Used Memory (KB): 284520
Total Memory (KB): 417600
Free Memory (KB): 133079
Query String:
Return Status Code: 200
CPU Time (ms): 0
JDBC Queries Run: 0
JDBC Total Time: 0
JDBC Total Execution Time: 0
JDBC Total Row Count: 0
AMF Request: None
Bytes Sent: 0
Time to First Byte (ms): 0
Time to Last Byte (ms): 0 [Note: Data is still streaming]
Time to Stream Open (ms): 1
Time to Stream Close (ms): - [Note: Data is still streaming]
User Agent: Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/535.19 (KHTML, like Gecko) Chrome/18.0.1025.168 Safari/535.19

This is for notification only. No special action was taken on the request.
Running Requests:
5 1337071199613 (2012-05-15 10:39:59.613) 1304ms 127.0.0.1 http://127.0.0.1:8500/cf/sleep.cfm
Complete details for all running requests appear after the stack trace below.
No new alerts will be generated until Tue

```
May 15 10:41:00 CEST 2012
Alert complete.
FusionReactor

JVM Stack Trace
---------------
JVM ID: 2
Thread ID: Reference Handler
Priority: 10
Hashcode: 60543300
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.Object.wait(Object.java:485)
java.lang.ref.Reference$ReferenceHandler.r
un(Reference.java:116)
----------------------------------------
----------------------------------------
JVM ID: 3
Thread ID: Finalizer
Priority: 8
Hashcode: 1142635176
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:118)
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:134)
java.lang.ref.Finalizer$FinalizerThread.ru
n(Finalizer.java:159)
----------------------------------------
----------------------------------------

Running Requests (Full Details)
```

```
--------------------------------
Request ID: 5
 Request URL:
http://127.0.0.1:8500/cf/sleep.cfm
 Status: EXECUTING
 Started (Milliseconds): 1337071199613
 Started (Date/Time): 2012-05-15
10:39:59.613
 Thread ID: web-0
 Client IP Address: 127.0.0.1
 Request Method: GET
 Execution Time (ms): 1304
 Used Memory (percentage): 61
 Max Memory (KB): 466048
 Used Memory (KB): 284520
 Total Memory (KB): 417600
 Free Memory (KB): 133079
 Query String:
 Return Status Code: 200
 CPU Time (ms): 0
 JDBC Queries Run: 0
 JDBC Total Time: 0
 JDBC Total Execution Time: 0
 JDBC Total Row Count: 0
 AMF Request: None
 Bytes Sent: 0
 Time to First Byte (ms): 0
 Time to Last Byte (ms): 0 [Note: Data is
still streaming]
 Time to Stream Open (ms): 1
 Time to Stream Close (ms): - [Note: Data
is still streaming]
```

User Agent: Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/535.19 (KHTML, like Gecko) Chrome/18.0.1025.168 Safari/535.19

[Memory Protection] FusionReactor Crash Protection Alert [HP-Pdv6-Laptop]

Hello, this is the FusionReactor Crash Protection agent at HP-Pdv6-Laptop, instance: coldfusion.cfmx9.HP-Pdv6-Laptop.
http://HP-Pdv6-Laptop:8088/fusionreactor/fhtml.cfm?mode=main
The following alert was generated at Tue May 15 11:07:27 CEST 2012.
[Rejected] Crash Protection - Memory Shortage Alert: Actual value: 35% (163MB), Threshold value: 100% (455MB)
Current (Triggering) Request:
Request ID: 39
Request URL: http://127.0.0.1:8500/cf/tightloop.cfm
Status: PENDING
Started (Milliseconds): 1337072847614
Started (Date/Time): 2012-05-15 11:07:27.614
Thread ID: web-5
Client IP Address: 127.0.0.1
Request Method: GET
Execution Time (ms): 1
Used Memory (percentage): 64
Max Memory (KB): 466048

Used Memory (KB): 298498
Total Memory (KB): 396928
Free Memory (KB): 98429
Query String:
Return Status Code: 200
CPU Time (ms): 0
JDBC Queries Run: 0
JDBC Total Time: 0
JDBC Total Execution Time: 0
JDBC Total Row Count: 0
AMF Request: None
Bytes Sent: 0
Time to First Byte (ms): 0
Time to Last Byte (ms): 0 [Note: Data is still streaming]
Time to Stream Open (ms): 0
Time to Stream Close (ms): - [Note: Data is still streaming]
User Agent: Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/535.19 (KHTML, like Gecko) Chrome/18.0.1025.168 Safari/535.19

This request was immediately rejected and did not run.
Running Requests:
None
****Complete details for all running requests appear after the stack trace below.****
No new alerts will be generated until Tue May 15 11:08:27 CEST 2012
Alert complete.
FusionReactor

```
JVM Stack Trace
---------------
JVM ID: 2
Thread ID: Reference Handler
Priority: 10
Hashcode: 60543300
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.Object.wait(Object.java:485)
java.lang.ref.Reference$ReferenceHandler.r
un(Reference.java:116)
-----------------------------------------
-----------------------------------
JVM ID: 3
Thread ID: Finalizer
Priority: 8
Hashcode: 1142635176
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:118)
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:134)
java.lang.ref.Finalizer$FinalizerThread.ru
n(Finalizer.java:159)
-----------------------------------------
-----------------------------------
JVM ID: 13
Thread ID: RMI TCP Accept-0
Priority: 5
Hashcode: 513860152
java.net.PlainSocketImpl.socketAccept(Plai
```

```
nSocketImpl.java:???)[Native Method]
java.net.PlainSocketImpl.accept(PlainSocke
tImpl.java:390)
java.net.ServerSocket.implAccept(ServerSoc
ket.java:453)
java.net.ServerSocket.accept(ServerSocket.
java:421)
sun.rmi.transport.tcp.TCPTransport$AcceptL
oop.executeAcceptLoop(TCPTransport.java:36
9)
sun.rmi.transport.tcp.TCPTransport$AcceptL
oop.run(TCPTransport.java:341)
java.lang.Thread.run(Thread.java:619)
----------------------------------------
----------------------------------------
JVM ID: 14
Thread ID: RMI Reaper
Priority: 5
Hashcode: 1535770241
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:118)
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:134)
sun.rmi.transport.ObjectTable$Reaper.run(O
bjectTable.java:333)
java.lang.Thread.run(Thread.java:619)
----------------------------------------
----------------------------------------
JVM ID: 15
Thread ID: GC Daemon
Priority: 2
```

```
Hashcode: 1650016544
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
sun.misc.GC$Daemon.run(GC.java:100)
----------------------------------------
----------------------------------
JVM ID: 22
Thread ID: RMI Scheduler(0)
Priority: 5
Hashcode: 1282888134
sun.misc.Unsafe.park(Unsafe.java:???)[Nati
ve Method]
java.util.concurrent.locks.LockSupport.par
k(LockSupport.java:158)
java.util.concurrent.locks.AbstractQueuedS
ynchronizer$ConditionObject.await(Abstract
QueuedSynchronizer.java:1925)
java.util.concurrent.DelayQueue.take(Delay
Queue.java:160)
java.util.concurrent.ScheduledThreadPoolEx
ecutor$DelayedWorkQueue.take(ScheduledThre
adPoolExecutor.java:583)
java.util.concurrent.ScheduledThreadPoolEx
ecutor$DelayedWorkQueue.take(ScheduledThre
adPoolExecutor.java:576)
java.util.concurrent.ThreadPoolExecutor.ge
tTask(ThreadPoolExecutor.java:947)
java.util.concurrent.ThreadPoolExecutor$Wo
rker.run(ThreadPoolExecutor.java:907)
java.lang.Thread.run(Thread.java:619)
----------------------------------------
----------------------------------
JVM ID: 118
```

```
Thread ID: Java2D Disposer
Priority: 10
Hashcode: 1785106435
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:118)
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:134)
sun.java2d.Disposer.run(Disposer.java:125)
java.lang.Thread.run(Thread.java:619)
----------------------------------------
----------------------------------------
```

```
Running Requests (Full Details)
-------------------------------
None
```

[Request Protection] FusionReactor Crash Protection Alert [HP-Pdv6-Laptop]

```
Hello, this is the FusionReactor Crash
Protection agent at HP-Pdv6-Laptop,
instance: coldfusion.cfmx9.HP-Pdv6-Laptop.
http://HP-Pdv6-Laptop:8088/fusionreactor/f
html.cfm?mode=main
The following alert was generated at Tue
May 15 10:47:35 CEST 2012.
[Notify] Crash Protection - Running
Request Quantity Alert: Actual value: 1
requests, threshold value: 1 requests
Current (Triggering) Request:
Request ID: 21
 Request URL:
http://127.0.0.1:8500/cf/sleep.cfm
 Status: PENDING
 Started (Milliseconds): 1337071655942
 Started (Date/Time): 2012-05-15
10:47:35.942
 Thread ID: web-5
 Client IP Address: 127.0.0.1
 Request Method: GET
 Execution Time (ms): 1
 Used Memory (percentage): 65
 Max Memory (KB): 466048
 Used Memory (KB): 306378
 Total Memory (KB): 396480
```

Free Memory (KB): 90101
Query String:
Return Status Code: 200
CPU Time (ms): 0
JDBC Queries Run: 0
JDBC Total Time: 0
JDBC Total Execution Time: 0
JDBC Total Row Count: 0
AMF Request: None
Bytes Sent: 0
Time to First Byte (ms): 0
Time to Last Byte (ms): 0 [Note: Data is still streaming]
Time to Stream Open (ms): 0
Time to Stream Close (ms): - [Note: Data is still streaming]
User Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:11.0) Gecko/20100101 Firefox/11.0

This is for notification only. No special action was taken on the request.
Running Requests:
18 1337071445943 (2012-05-15 10:44:05.943) 210000ms 127.0.0.1 http://127.0.0.1:8500/cf/sleep.cfm
****Complete details for all running requests appear after the stack trace below.****
No new alerts will be generated until Tue May 15 10:48:35 CEST 2012
Alert complete.
FusionReactor

```
JVM Stack Trace
---------------
JVM ID: 2
Thread ID: Reference Handler
Priority: 10
Hashcode: 60543300
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.Object.wait(Object.java:485)
java.lang.ref.Reference$ReferenceHandler.r
un(Reference.java:116)
-----------------------------------------
-----------------------------------
JVM ID: 3
Thread ID: Finalizer
Priority: 8
Hashcode: 1142635176
java.lang.Object.wait(Object.java:???)[Nat
ive Method]
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:118)
java.lang.ref.ReferenceQueue.remove(Refere
nceQueue.java:134)
java.lang.ref.Finalizer$FinalizerThread.ru
n(Finalizer.java:159)
-----------------------------------------
-----------------------------------
Running Requests (Full Details)
-------------------------------
Request ID: 18
 Request URL:
http://127.0.0.1:8500/cf/sleep.cfm
```

Status: EXECUTING
Started (Milliseconds): 1337071445943
Started (Date/Time): 2012-05-15 10:44:05.943
Thread ID: web-4
Client IP Address: 127.0.0.1
Request Method: GET
Execution Time (ms): 210000
Used Memory (percentage): 62
Max Memory (KB): 466048
Used Memory (KB): 291765
Total Memory (KB): 394304
Free Memory (KB): 102538
Query String:
Return Status Code: 200
CPU Time (ms): 0
JDBC Queries Run: 0
JDBC Total Time: 0
JDBC Total Execution Time: 0
JDBC Total Row Count: 0
AMF Request: None
Bytes Sent: 0
Time to First Byte (ms): 0
Time to Last Byte (ms): 0 [Note: Data is still streaming]
Time to Stream Open (ms): 1
Time to Stream Close (ms): - [Note: Data is still streaming]

```
User Agent: Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/535.19 (KHTML, like Gecko) Chrome/18.0.1025.168 Safari/535.19
```

Using the FusionReactor JDBC Driver Wrapper

Introducing FusionReactor JDBC Driver Wrapper

This documentation will take you through the steps to install the FusionReactor JDBC Driver Wrapper, and provide you with a few examples of how it can be used.

The intention of the Wrapper is to provide a thin layer between J2EE applications and JDBC-accessible databases, in order to intercept and observe the interaction between them.

The Wrapper can then report various useful metrics to FusionReactor, and can also step in to prevent runaway queries. The Wrapper is fully integrated into the FusionReactor Administrator, and you can easily see all the queries which ran during a request, together with their runtime metrics and row counts.

Although the Wrapper is equally useful for JSP and Servlet deployment, this tutorial will illustrate the usage of the Wrapper under ColdFusion.

> **NOTE**
>
> The Wrapper has its own User Guide, which comprehensively details all possible options, together with the JDBC URL syntax and error messages.
>
> You should consult the User Guide when working through the tutorial. The JDBC Driver Wrapper User Guide is the definitive reference for the driver.

Basic Setup of the JDBC Wrapper

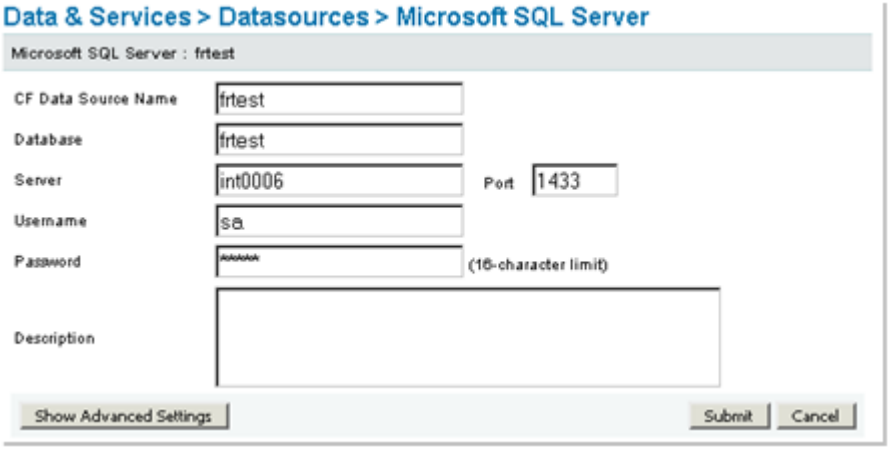
Gathering metrics from ColdFusion pages

Pathologically bad pages

Basic Setup of the JDBC Wrapper

The first step in using the Wrapper is to find your ColdFusion datasource. This datasource should be working prior to applying the Wrapper. Find your datasource within the ColdFusion Administrator's 'Data Sources'

section. Here's ours:



You can see that the example will use a built-in Macromedia driver to connect to the '**frtest**' SQL Server database on **int0006.**

In order to have FusionReactor gather metrics for this connection, we must 'wrap' the driver. This will involve converting it to an 'other' type driver.

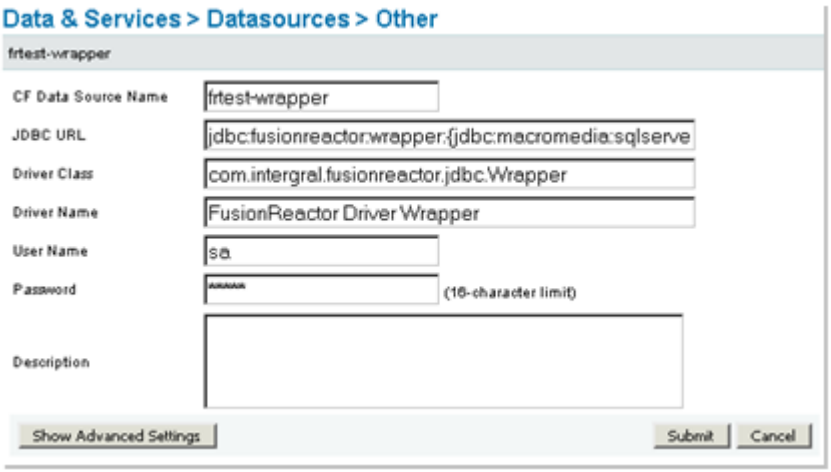
In order to wrap the driver, you'll need to obtain the JDBC URL for your existing datasource. The easiest way to do this is to click on **Server Settings -> Settings Summary**, then scroll down until you find the existing datasource. Copy the JDBC URL.



frtest	
CF data source name	frtest
Description	
Driver	MSSQLServer
JDBC URL	jdbc:macromedia:sqlserver://int0006:1433;databaseName=f
Username	sa
Login timeout	30 seconds
Long text buffer size	64000
Timeout	1200 seconds
Maintain connections	Yes
Interval	420 seconds
Restricted SQL operations	
Disable connections	No

In ColdFusion Administrator, add a new driver called 'frtest-wrapper', of type 'other'. In the **JDBC URL** field, you will enter the JDBC URL you copied from the Settings Summary page earlier, wrapped using FusionReactor's JDBC Wrapper syntax (more details below).
You can refer to the FusionReactor JDBC Driver Wrapper User Guide to get some examples of JDBC URLs for the Macromedia builtin drivers.

Here's the form with our previous connection 'wrapped' with the FusionReactor JDBC Driver Wrapper.



The important things to note here are:

1. Since we're using an 'other' type driver, we used the JDBC URL we copied out of the Settings Summary page.

```
jdbc:macromedia:sqlserver://int0006:1433;databaseName=frtest;SelectMethod=direct;sendStringParametersAsUnicode=false;MaxPooledStatements=1000
```

We've added the FusionReactor JDBC Driver Wrapper text to the JDBC URL:

```
jdbc:fusionreactor:wrapper:{
jdbc:macromedia:... URL }
```

2. We've add the FusionReactor Driver Wrapper class name:

```
com.intergral.fusionreactor.jdbc.Wrapper
```

3. The username and password are unchanged.

Since we're using a builtin Macromedia driver, there's no need to supply the 'driver' argument to the FusionReactor JDBC URL – CF will pre-load this driver for us. If we were using a third party driver, we would simply add the classname to the JDBC URL, like this:

```
jdbc:fusionreactor:wrapper:{
jdbc:mysql:... URL
};driver=com.somecompany.Driver
```

When you submit this form, ColdFusion will load our driver and make sure it can access the wrapped database.



If all goes well, you'll see an 'OK' in the status column: If something goes wrong, you'll get plenty of information to help you fix it. If FusionReactor can't access the wrapped driver, it will give you a helpful message indicating where the problem lies. If the wrapped driver has problems accessing the database, FusionReactor will relay its message so you can fix the problem. Here's an example of a problem – we got the name of the FusionReactor driver class wrong:



You can see that ColdFusion couldn't load a suitable driver. Here's an example showing an incorrect database username:



Simply click on the datasource name to go back to the form and fix the problems. Once the datasource is ready, you can begin using it to gather metrics.

```
<cfquery datasource="frtest-wrapper">
      select * from usr
</cfquery>

<cfoutput>Read #getusers.recordCount#
users</cfoutput>
```

Gathering metrics from ColdFusion pages

Once you've set up the frtest-wrapper datasource and verified it, we can begin to explore the metric-gathering

abilities of FusionReactor.

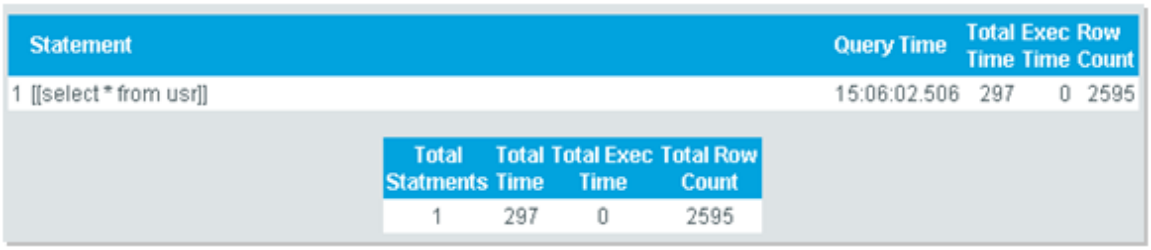
Using our 'frtest' database, we can easily write a ColdFusion query using our new datasource:

As you can see, this code simply reads all the rows from our 'users' table, a fairly small dataset. Running this page (which I've called 'tutorial1.cfm') gives the following output:

If we go to FusionReactor's 'Request History' screen, we can see this request as the most recent.



Clicking on the blue 'data' icon at the very left of the request brings up the 'Request Details' page. When that page has displayed, click on the 'JDBC' tab heading to bring up the JDBC Details page, which looks something like this:



Statement	Query Time	Total Time	Exec Time	Row Count
1 [[select * from usr]]	15:06:02.506	297	0	2595

Total Statments	Total Time	Total Exec Time	Total Row Count
1	297	0	2595

This page tells us the query text that ran (if you were using CFQUERYPARAM, the actual, real values would be displayed in the statement), the exact time the statement ran, the total time (from running the query to finishing with the data), the exec time (the actual time the database took to run the query – in this case 0ms) and the number of rows read out by your query.

The final block contains the totals for all the queries which ran on the page.

Why did the query run in 0ms?

You may occasionally see Total Time and Exec Time values which are 0ms. There are a couple of explanations for this.

1. The query never actually ran against the database.

If the query has been run before, and if there were no conditions which invalidated the underlying driver's cache, the result set may be delivered without actually running the query on the database. This can save a lot of time, and yields an Exec Time of 0ms.

2. The query ran, but ran very quickly.

In some operating systems, notably the Windows Server line (including Windows XP) the highest resolution of a timer available to Java is 10ms. If the query ran faster than 5ms, the timer will return a rounded-down value of 0ms. This is unfortunately unavoidable.

If you have more queries which run on a page, they'll be listed in the order they complete.

Also See

Basic Setup of the JDBC Wrapper

JDBC Driver Wrapper User Guide

Pathologically bad pages

Let's have a quick look at a couple of pathologically bad queries – designed to produce lots of data. Here's our

test query:

The engineer is trying to work out which user (u) owns which object (o) but has gone for his lunch and on his return, forgets to add a constraining 'where' clause.

This query is going to select a cross join (cartesian product) of all Users against all Objects. The user table, as we've already seen, contains 2595 rows, and the object table contains 24799 rows. This will yield a Cartesian product of 2595 x 24799 = 64,353,405 rows.

Oblivious to the error, we run the page and notice it's taking a long time to complete. In the FusionReactor Administrator, we can find the currently-running request:



... and click on the blue 'data' icon (top right of the icon block) to bring up the 'Request Details' page for this request. Once that page has appeared, we can click on the JDBC tab to bring up the JDBC Details page, which looks like this:



Currently Running Statement	Query Date	Query Time
[[select * from usr u, obj o]]	15:24:15.18	138734

It's now clear that our statement is causing the problem – since it's listed as the 'Currently Running Statement'. We can then use the red 'X' icon on the Running Requests page to kill this request. But what if we wanted to limit the number of rows retrieved to some 'hard' value, to prevent runaway queries in production?

Using the Row Limiter

The FusionReactor JDBC Driver Wrapper has a builtin Row Limiter, which is configured on the Data Sources page of the CF Administrator.

In the CF Administrator, edit the 'frtest-wrapper' data source, and append

```
;rowLimit=10000
```

to the JDBC URL, making it:

```
jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://int0006:1433;databaseName=frtest;SelectMethod=direct;sendStringParametersAsUnicode=false;MaxPooledStatements=1000};rowLimit=10000
```

Submitting that form will make sure the syntax is correct. The Driver Wrapper will now limit the number of returned rows to 10,000.

If we run our rogue page again, we can see that it only reads 10,000 rows -- at which point the Row Limiter activates to prevent the query running away:

Statement	Query Time	Total Time	Exec Time	Row Count
1 [[select * from usr u, obj o]]	15:36:39.403	3391	31	10000

... and the JDBC Details page for the request also reflects the limit

Using FusionReactor Enterprise Scripting

Introducing FusionReactor Enterprise Scripting

FusionReactor Enterprise Scripting (available as part of the FusionReactor Enterprise Edition), enhances FusionReactor's Enterprise Monitoring to trigger configurable scripts when a server's responsiveness status changes.

This feature can be used to perform any task which can be launched from a platform script. Examples might include:

- Interacting with SNMP or enterprise monitoring/reporting systems
- Performing automated restarts of affected instances
- Sending custom email messages
- Writing custom log messages

Since FusionReactor Enterprise Scripting is able to launch any platform executable (shell script, executable binary etc.), it may be tailored to virtually any environment.

Intended Audience

This technical document is targeted at FusionReactor administrators who are responsible for monitoring one or more FusionReactor servers using the FusionReactor Enterprise Dashboard. This document will guide the administrator through the new functionality, illustrating how and when FusionReactor runs scripts, and the best practices and caveats involved with automated scripting.

FusionReactor Enterprise Scripting in a Nutshell

- What Are Enterprise Scripts?
- How Do I Configure an Enterprise Script?
- When Does FusionReactor Run Enterprise Scripts?
- How Does FusionReactor Run These Scripts?
- Logging Scripts
- Operational Impacts of Scripting

Using the Example Scripts

FusionReactor Enterprise Scripting in a Nutshell

What Are Enterprise Scripts?

How Do I Configure an Enterprise Script?

When Does FusionReactor Run Enterprise Scripts?

How Does FusionReactor Run These Scripts?

Logging Scripts

Operational Impacts of Scripting

What Are Enterprise Scripts?

FusionReactor (Enterprise Edition) can trigger a script when a server becomes unresponsive, and when it becomes responsive again. This mechanism might allow you to pro-actively attempt to restart a failed server or instance, integrate FusionReactor into an existing monitoring environment or provide custom logging or reporting. Any program which is runnable on your FusionReactor monitoring system may be used an Enterprise Script target.

On **Unix** and **Unix-based** systems, this includes everything you can start from a command line, for instance:

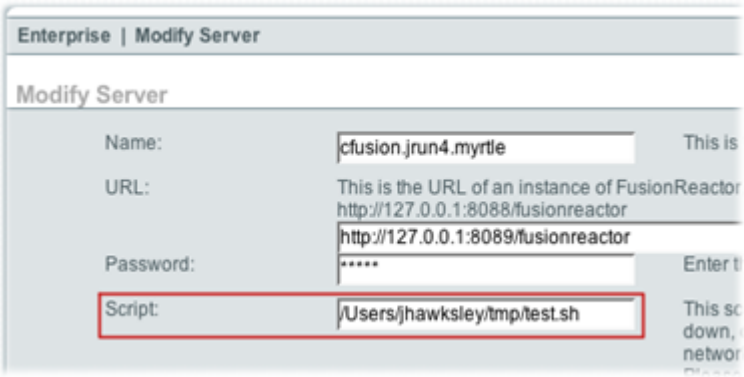
- Binary programs
- Shell scripts (including Bash, Ruby, Python and Perl)
- Java programs (when launched from an appropriate shell script)

On **Windows** systems, anything that is runnable as a binary, or can be run from the command prompt, for instance:

- Binary programs
- Batch files

How Do I Configure an Enterprise Script?

Enterprise scripts are configured by editing the **Script** of the server's property page, accessible in the Enterprise -> Manage Servers page, then clicking on the **Modify** icon ( ) of the required server.



You should take care to ensure the full path and filename to the script are correct.

When Does FusionReactor Run Enterprise Scripts?

Enterprise Scripts are run whenever the Enterprise Monitor detects that a monitored instance has changed state:

- an instance which was previously available is no longer providing Enterprise data
- an instance which was previously unavailable has begun to provide Enterprise data

Scripts are run only if an instance changes state while it is being observed.

Additionally, scripts are only run if:

- The Enterprise Server Alerting system is running
- This is configured in **Enterprise Settings -> Server Shutdown/Start Up Alerts**
- At least "On Shutdown" must be selected. If you wish to run scripts when an instance becomes available, "On Shutdown and Start Up" must be selected.
- If you do not wish to additionally receive email for these events, disable notification in **FusionReactor -> Settings -> Email Server -> Notification.**

How Does FusionReactor Run These Scripts?

Launch Mechanism

FusionReactor runs these scripts by spawning them using Java system commands. The scripts will be run in the context of the user under which your J2EE (ColdFusion) server runs. This user must have at least **read + execute** access to these scripts. Any files or other executables called by the script must also be accessible by this user.

The script will be run with the **current working directory** (CWD) of the J2EE application server. Because of the variety of platforms available, this may be unpredictable. Any scripts you write should therefore **not** use the current directory notation (usually a single dot) to address files. If you plan to access files within the script, their paths should be specified completely.

Script Arguments

FusionReactor supplies several command-line arguments to the script. These arguments may be used by the script to perform logging or restart operations. The supplied arguments are (in order):

Argument	Description
UP or **DOWN**	Reflecting the instance status.
instance name	As registered in the **Manage Servers** screen.
IP Address	As returned from a DNS lookup of the machine name part of the URL used to monitor this instance.
Process ID	If available, the process ID of the J2EE application server on the remote machine. If the FusionReactor native library is not available, or FusionReactor could not read this value, this field will be **-1.**
Last Seen Time	The time, measured in milliseconds from midnight on January 1st 1970, which the server was last successfully polled for Enterprise data. If the server has not been observed as running during this session, this field will be **-1.**

Logging Scripts

When FusionReactor fires a script, an appropriate message is written to the Crash Protection log, located in FusionReactor/instance/<instance_name>/log/crashproteciton-0.log. This log is shared with other CrashProtection messages, and not all fields are used by Enterprise scripting.

Field Name	Description
Date	Formatted date
Time	Formatted time, to millisecond precision
Time (Epoch)	Absolute time, measured in milliseconds from midnight on January 1st 1970
ServerUP or **ServerDOWN**	The triggering state for this script invocation

One of: **RUNSCRIPT** **SCRIPTREADFAILED** **SCRIPTEXCEPTION**	 Script successfully called Script not found, or not readable Script found and readable, but an exception occurred during run
6 Unused Fields	
Server ID	Server ID of the affected instance, configured within Server Manager
Script	Full command and arguments launched by FusionReactor
Unused Fields	The remainder of the log string is unused.

For the exceptional cases **SCRIPTREADFAILED** and **SCRIPTEXCEPTION**, FusionReactor will log the message associated with the exception to the FusionReactor log.

Operational Impacts of Scripting

There are a few points which should be borne in mind when configuring scripting.

System Restarts and Self-Monitoring

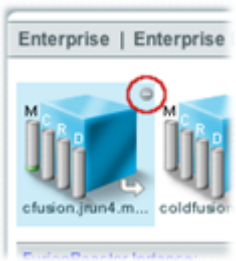
If FusionReactor is configured to monitor itself, i.e. is monitoring the same instance in which it is configured, scripting should **not** be used for operations which affect this instance. Because the order in which FusionReactor's subsystems start up, FusionReactor may prematurely fire a script while the 'self' instance is still starting up.

In these circumstances, we recommend transitioning your environment to a **High Availability** monitoring solution. This entails installing a new J2EE server (Tomcat, for instance), and installing FusionReactor into that. This container will be used purely as a FusionReactor host, and will be used to monitor other containers.

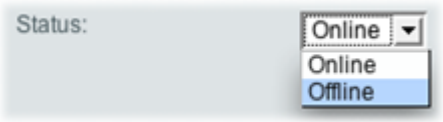
It may be necessary to create scripts which perform system reboots. Again, we recommend a careful evaluation of the impacts of this type of script before implementation. A script which restarts a system should not attempt to restart the system on which the monitoring solution runs.

Manual Restarts

If a script is configured for a given instance, it will be fired when that instance becomes unavailable. FusionReactor does not differentiate between overloaded (or failing) instances, and instances which have been deliberately stopped. Therefore, if you stop an instance manually, through Windows' Services panel for instance, FusionReactor will fire the configured script. As an operational matter, the affected instances should be **offlined** from FusionReactor **before** being shut down. This can be done within the Enterprise Dashboard, by clicking the **+/-** button on the server icon:



... or from within **Enterprise > Manage Servers** by selecting the **Modify** icon for the affected server, then changing its **Status** to **Offline**.



FusionReactor will not monitor these systems. When maintenance is complete, the servers should be onlined again by reversing the process.

Using the Example Scripts

Install the Scripts

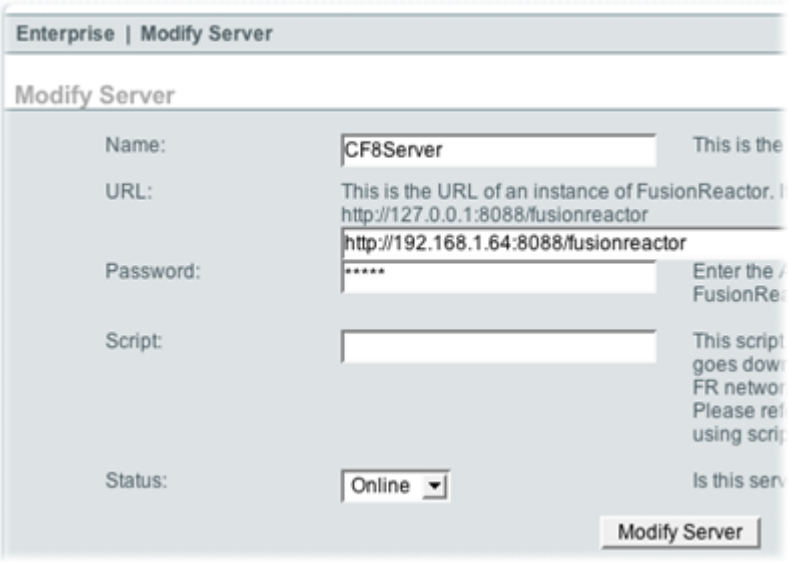
The example scripts are provided in /FusionReactor/etc/cp/, thereafter the structure is split into scripts which will run on Unix platforms, and those which will run on Windows platforms. You are free to run these scripts from this location, but we would recommend you copy these templates before editing them. You will then always have a pristine copy available for new scripts.

Worked Example Controlling Windows ColdFusion 8 from Windows

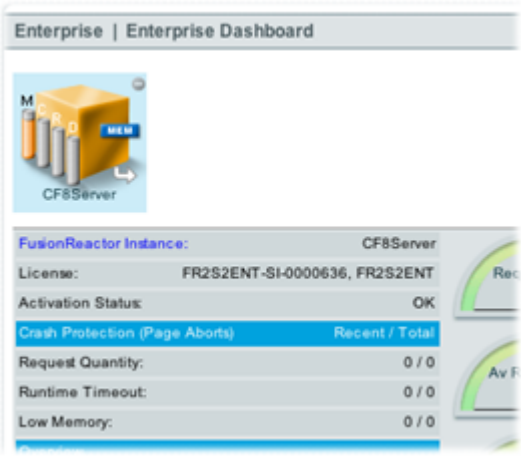
In order to get you started, we've provided you with a worked example. In our example scenario, we will use a FusionReactor Enterprise Edition instance on a ColdFusionMX 7 instance to monitor a ColdFusion 8 installation, also on Windows, also running FusionReactor Enterprise Edition.

Enterprise Dashboard

The first stage in preparing the environments is to ensure that both systems are running smoothly, and the monitor is able to poll the target system for enterprise data. We add the remote system to the monitor:



... and check that Enterprise Dashboard is retrieving information from it:



Script Preparation

For this example, we'll be using the restart-Coldfusion8-OnWindows.bat script from the FusionReactor/etc/cp/windows folder.

For our example, we copy the example script to a temporary folder, from where we can work on it:

copy restart-Coldfusion8-OnWindows.bat c:\tmp

In order to customize the script, we open it in an editor. All provided example scripts are commented extensively.

There are a couple of variables we must customize in the script:

- We set the LOGFILE (line 43) to c:\tmp\script.log
- We change the USER and PWD (lines 53 and 54 respectively) to reflect the Windows user with permissions to restart ColdFusion.

Adding the Script to Manage Servers

The final step in the configuration is to add the script to the monitored server's configuration. We edit the server's Enterprise Dashboard configuraiton by clicking on **Manage Servers,** then clicking the **edit** icon of the monitored server. We enter the script location in the **Script** field.



Testing the Script

The script can be tested by simply using the Windows Service control panel to stop the monitored ColdFusion 8 service. Observing the script log file c:\tmp\script.log file shows the output of the script. The ColdFusion 8 service can be observed restarting in the control panel.



Conclusion

We've shown you how to configure Enterprise Scripting to restart a ColdFusion 8 server. The scope for what scripts can do is immense, since there are no restrictions on what they may call. It would be a simple task, for example, to integrate FusionReactor into an SNMP monitoring solution, write custom log messages or send SMS text messages.

FusionReactor 4.5.x Installation Guide

Installing FusionReactor on Windows

- Installing FusionReactor 4.5.x on Windows
- Updating FusionReactor 3.x on Windows
- Updating FusionReactor 2.0 on Windows
- Updating FusionReactor 1.0 on Windows
- Console Installation on Windows
- Updating FusionReactor 4.0.x on Windows

Uninstalling FusionReactor on Windows

Installing FusionReactor on Linux

- Installing FusionReactor 4.5.x on Linux
- Updating FusionReactor 3.x on Linux
- Updating FusionReactor 2.0 on Linux
- Updating FusionReactor 1.0 on Linux
- Console Installation on Linux
- Updating FusionReactor 4.0.x on Linux

Uninstalling FusionReactor on Linux

Overview

Installing FusionReactor on Solaris

- Installing FusionReactor 4.5.x on Solaris
- Updating FusionReactor 3.x on Solaris
- Updating FusionReactor 2.0 on Solaris
- Updating FusionReactor 1.0 on Solaris
- Console Installation on Solaris

Installing FusionReactor on Mac

- Installing FusionReactor 4.5.x on Mac
- Updating FusionReactor 3.x on Mac

Uninstalling FusionReactor on Mac

Uninstalling FusionReactor on Solaris

Where are FusionReactor 4's Manual Installer Packages?

- Manual Instance Installation

Installing FusionReactor in Locked Down Environments

Installing FusionReactor on Windows

Installing FusionReactor 4.5.x on Windows

Updating FusionReactor 3.x on Windows

Updating FusionReactor 2.0 on Windows

Updating FusionReactor 1.0 on Windows

Console Installation on Windows

Updating FusionReactor 4.0.x on Windows

Installing FusionReactor 4.5.x on Windows

Executables

Executable	System Architecture
FusionReactor_windows_4_5_0.exe	32-bit
FusionReactor_windows_4_5_0_x64.exe	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

To start the installation, login as user *Administrator* and run the appropriate executable as described in the above section **Executables.**



Figure 1: Welcome screen

Figure 1: The first step in the installation wizard is the welcome screen.

Clicking the Next button will proceed with the installation and open the License Agreement as shown in Figure 2. Clicking Cancel will exit the Setup.

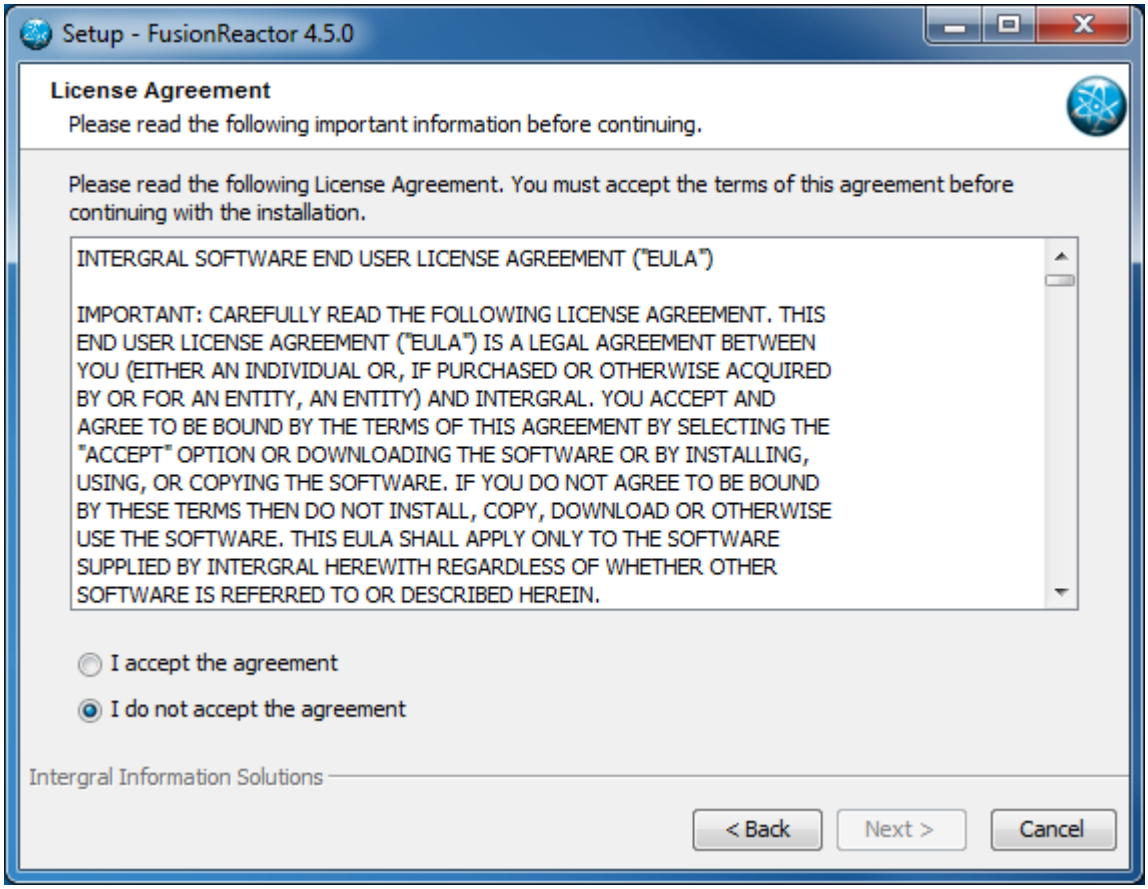


Figure 2: License Agreement

Figure 2: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.

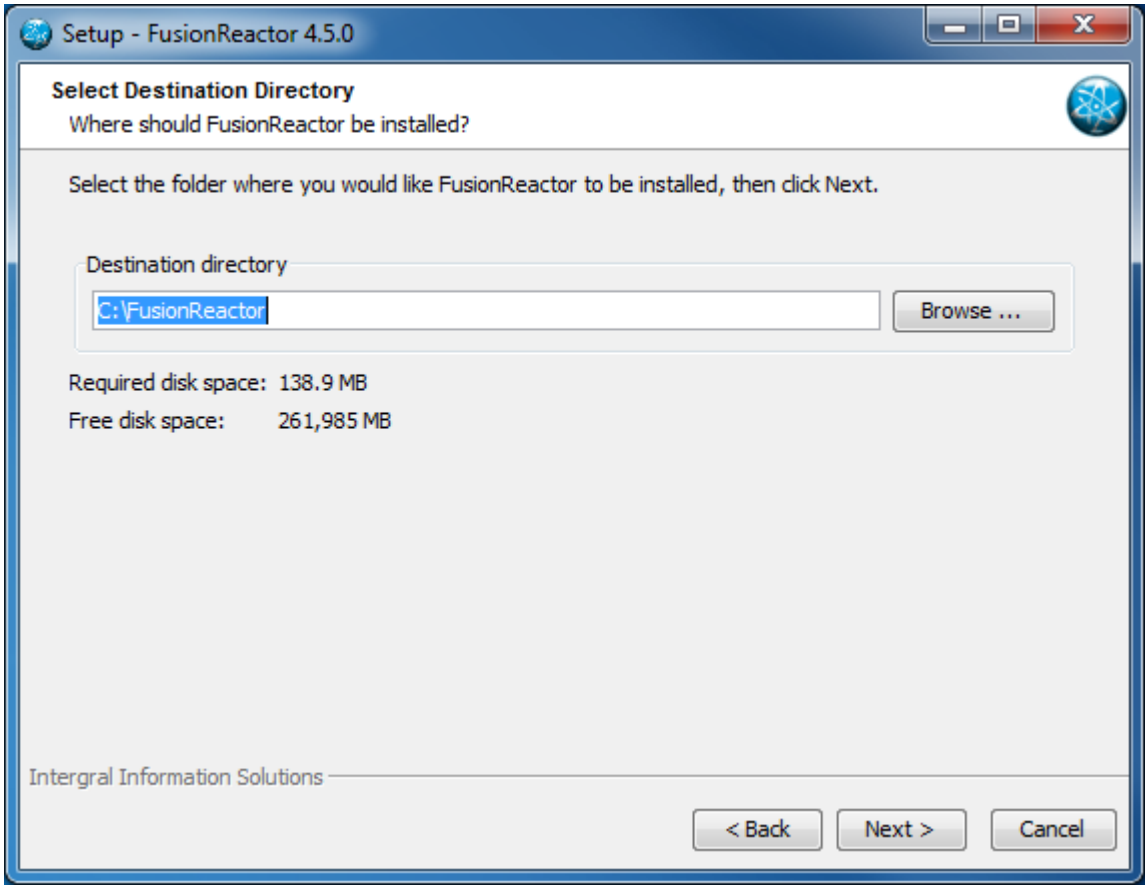


Figure 3: Installation Directory Selection

Figure 3: Now you have to select the directory in which FusionReactor should be installed. By default FusionReactor will be installed into `C:\FusionReactor\`

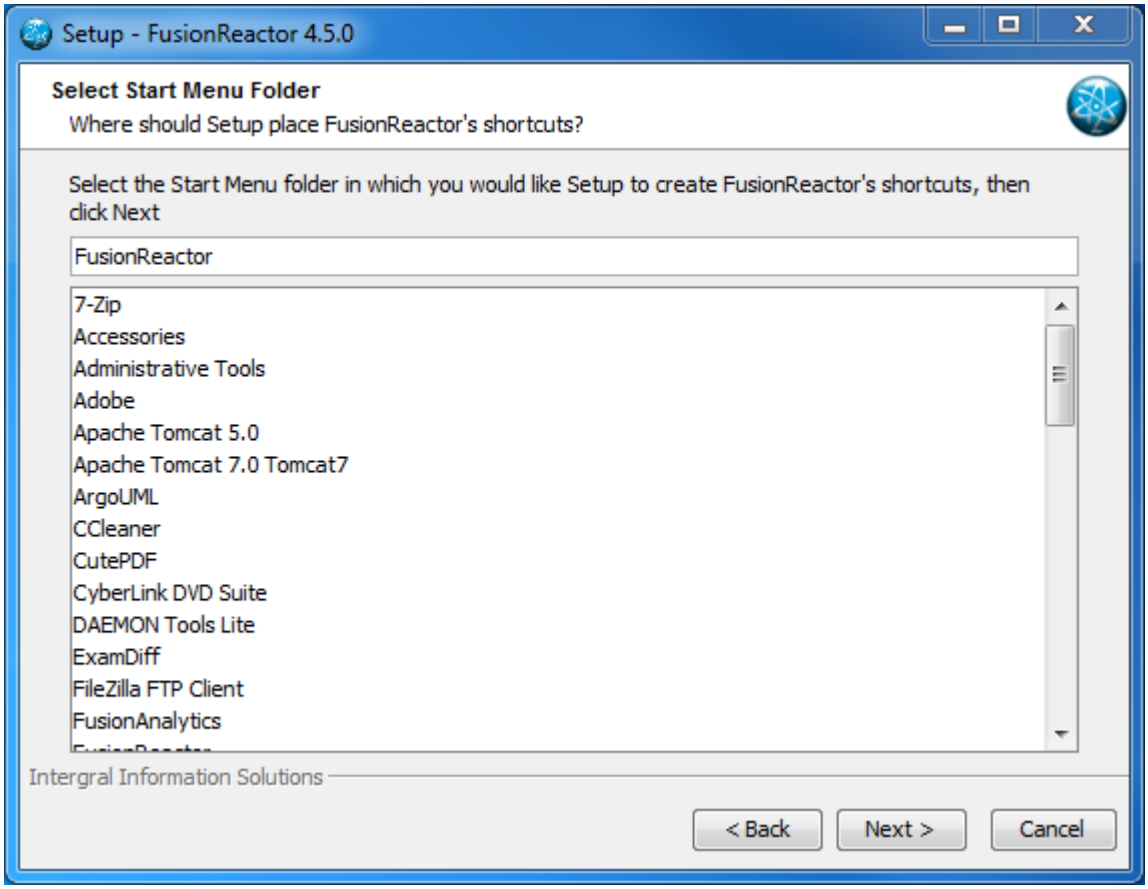


Figure 4: Start Menu Folder

Figure 4: Select the Start Menu folder in which FusionReactor shortcuts will be placed.



Figure 5: FusionReactor Administration Manager (FRAM)

Figure 5: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the server ports (See Figure 6). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 7).

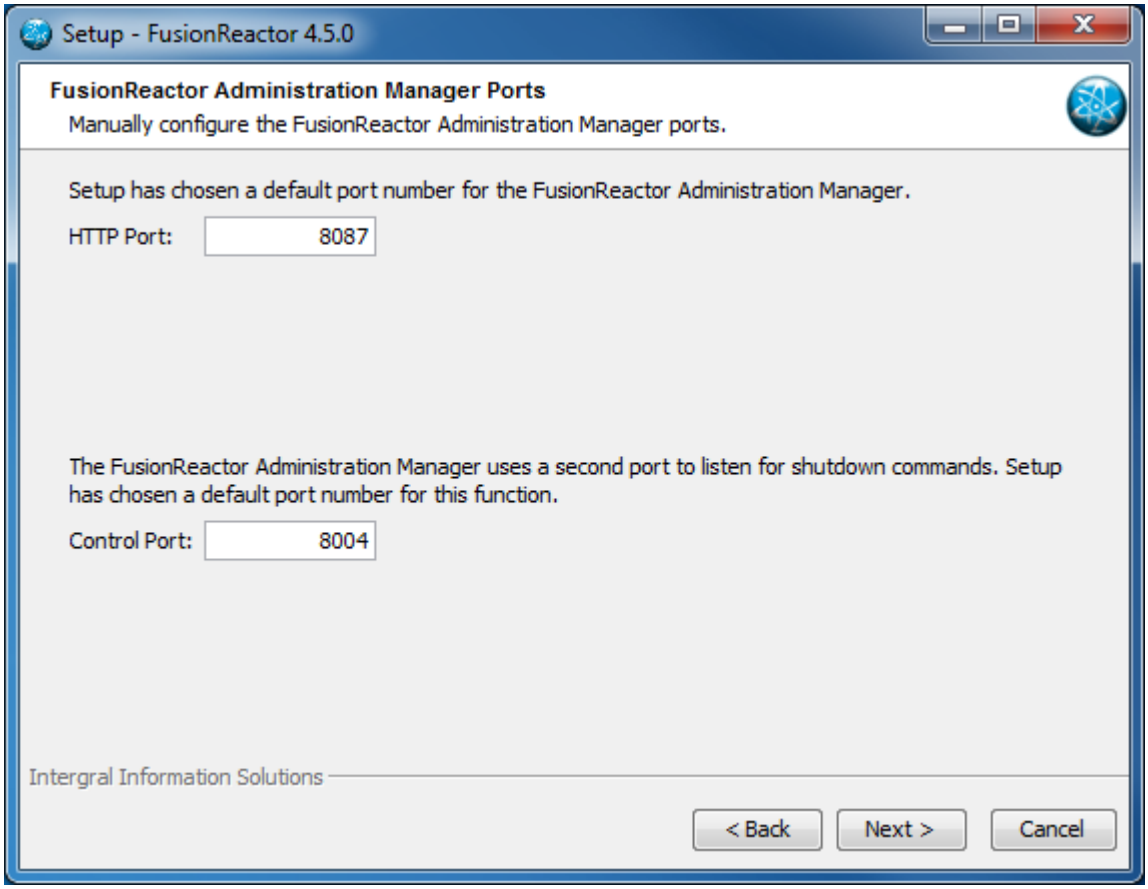


Figure 6: Manually Configuring Ports

Figure 6: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

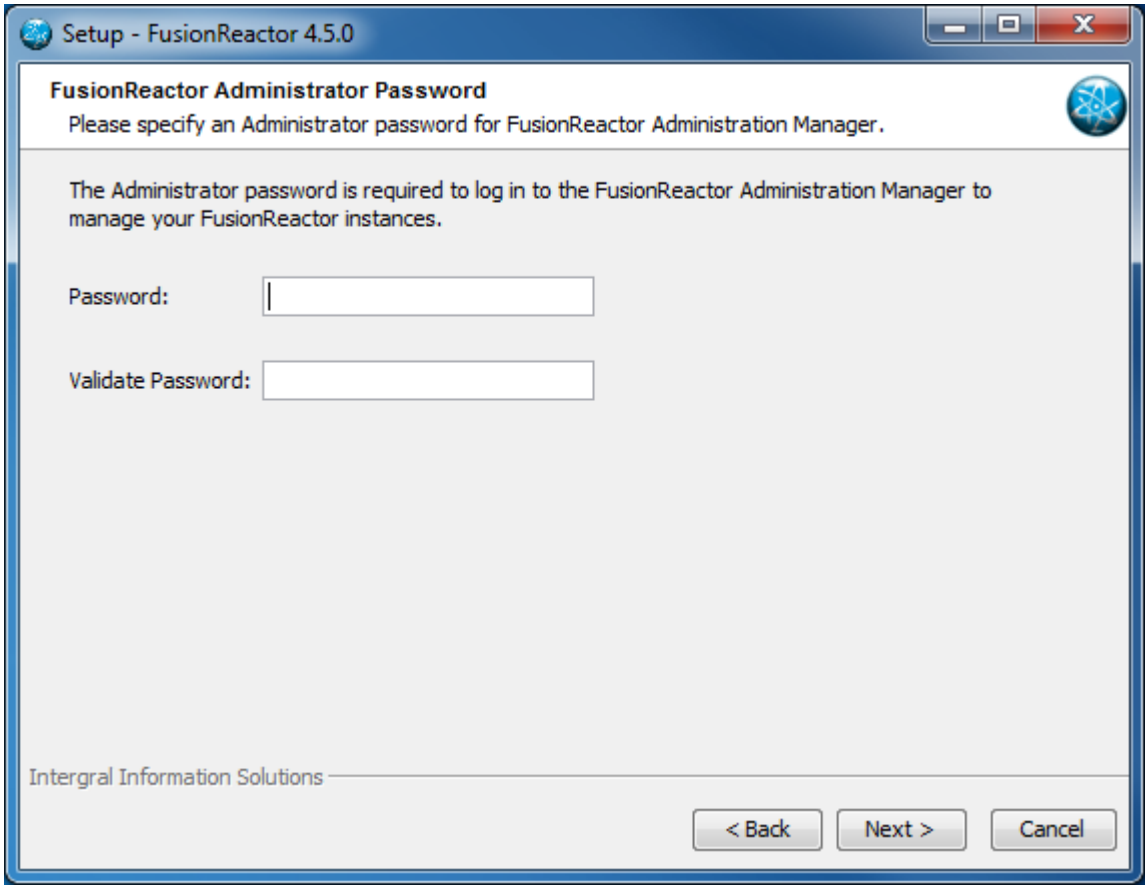


Figure 7: Administration Manager Password

Figure 7: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

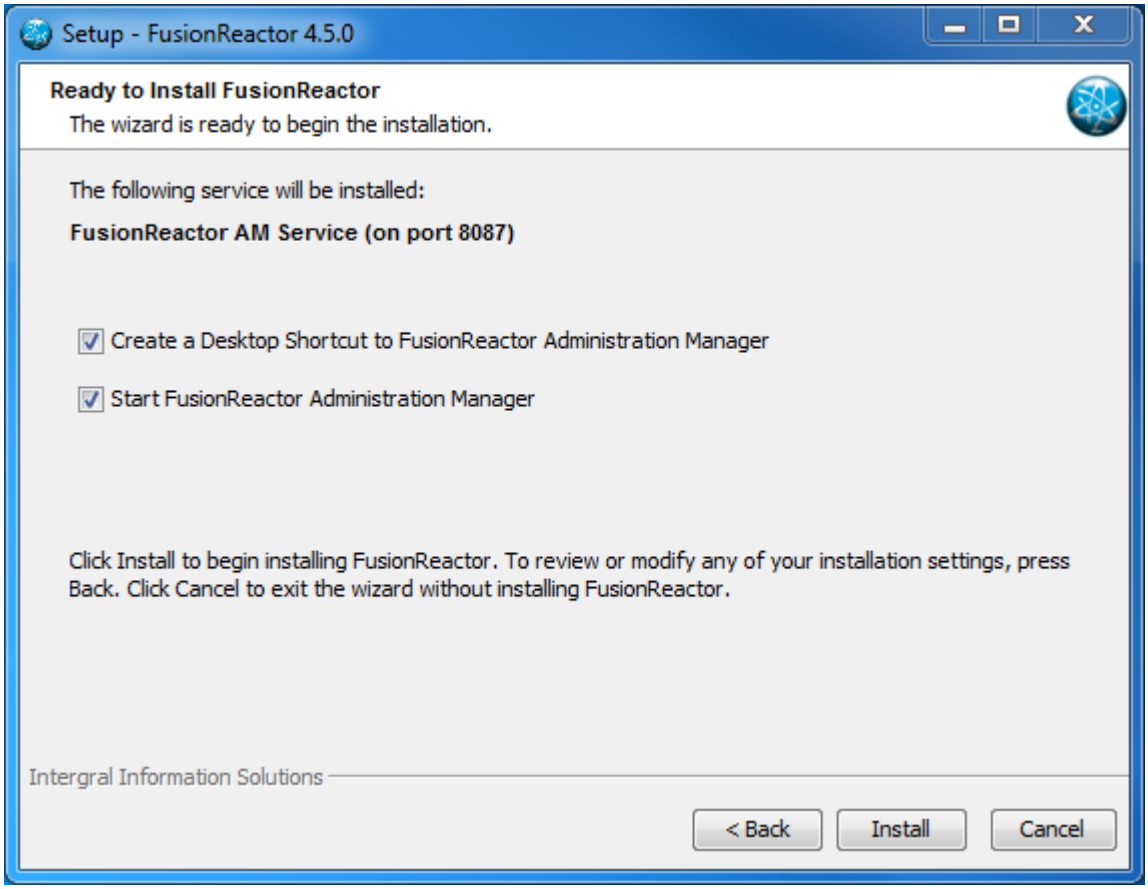


Figure 8: Ready to Install

Figure 8: The FusionReactor Setup can create a shortcut on the Desktop to FRAM, if you do not want Setup to perform this action uncheck the option *'Create Desktop Shortcut to FusionReactor Administration Manager".*

Setup can also automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin installing FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

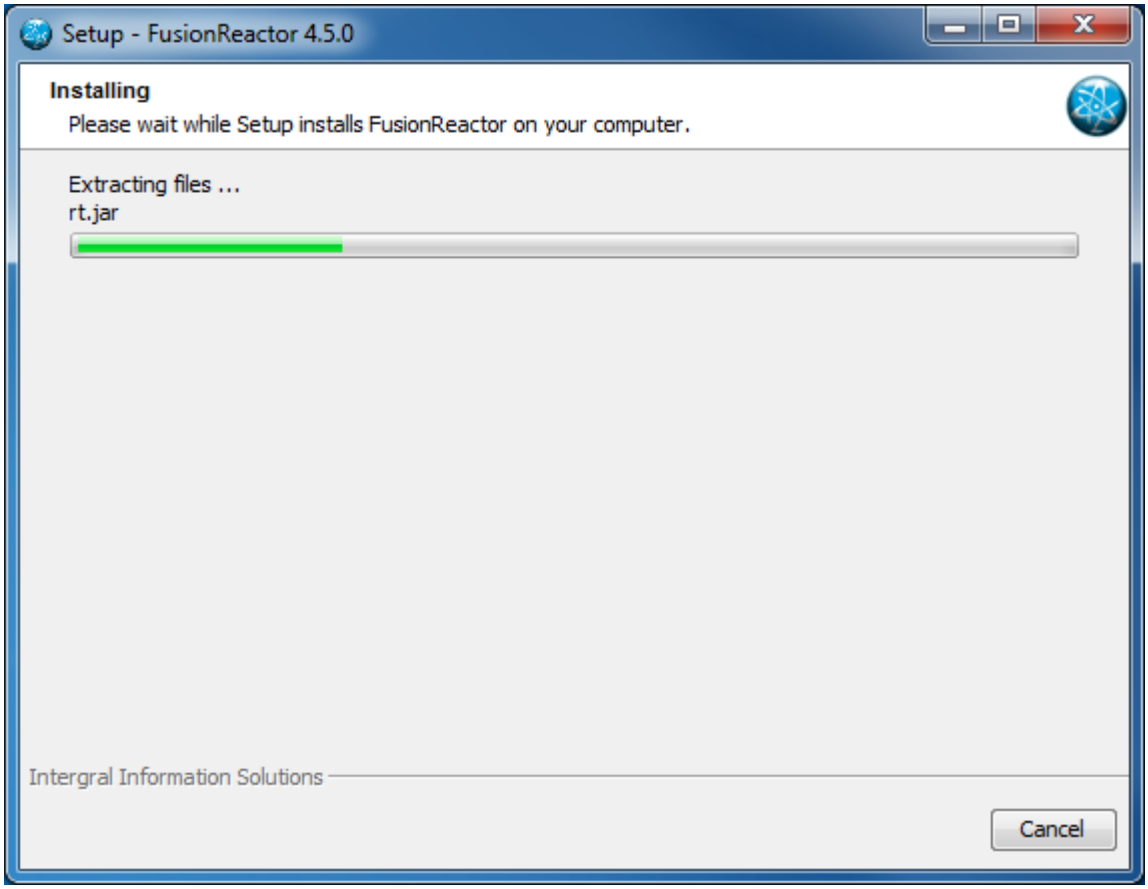


Figure 9: Installation Progress

Figure 9: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 10: Configure Instances Screen

Figure 10: After the installation has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. To configure now launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Server Scan Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close''* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 8.



Figure 11: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 11: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 8. This is the last screen in the setup, clicking Finish will close the window.

Updating FusionReactor 3.x on Windows

Executables

Executable	System Architecture
FusionReactor_windows_4_5_0.exe	32-bit
FusionReactor_windows_4_5_0_x64.exe	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.

Setup

You can run the FusionReactor 4.5.x Setup to update an already installed version of FusionReactor 3.x.

To start the update, login as user *Administrator* and run the appropriate executable as described in the above section **Executables.**



Figure 1: Welcome Screen

Figure 1: The first step in the installation wizard is the welcome screen as shown in Figure 1. Clicking the Next button will proceed with the update.

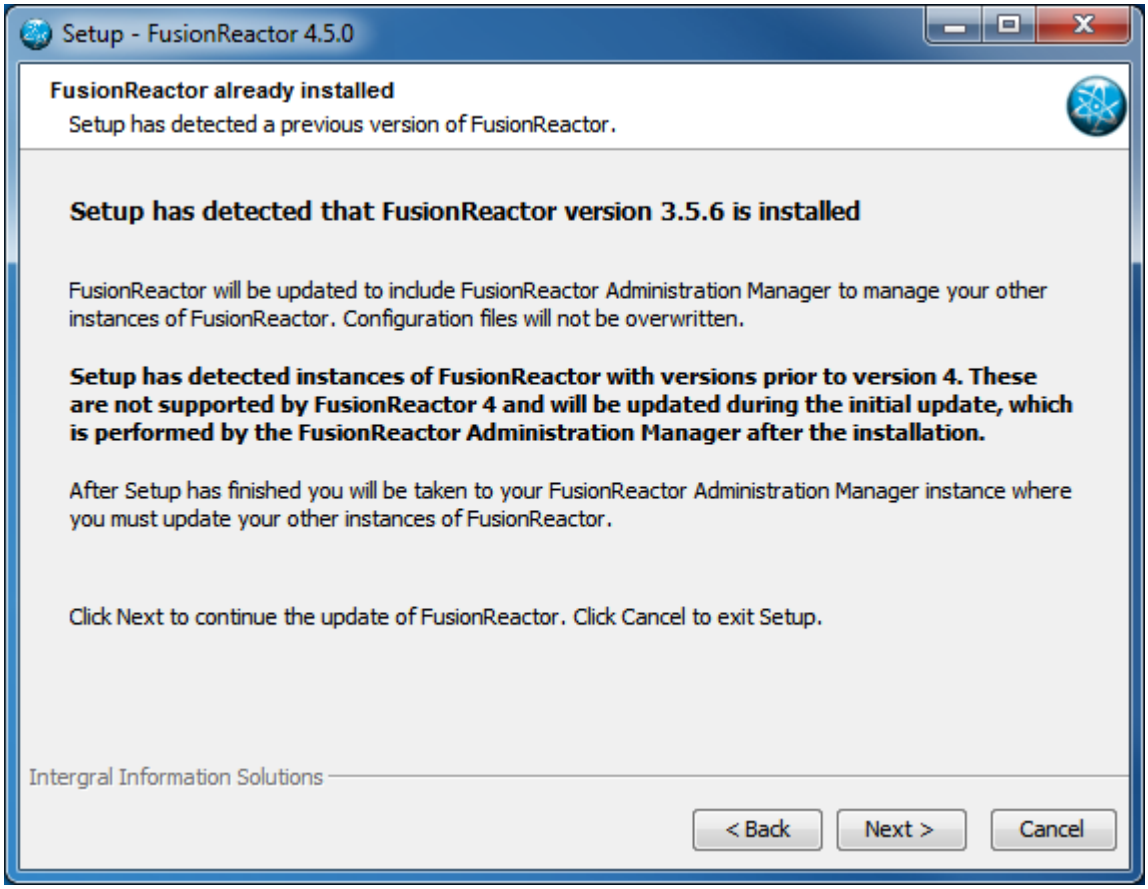


Figure 2: Update the Current Version

Figure 2: If there is already a version of FusionReactor 3.x on your computer that has been installed with the FusionReactor 3.x Setup, you will see a screen similar to the one shown. Carefully read the information and click Next to proceed with the update. Click Cancel to exit the installation wizard.

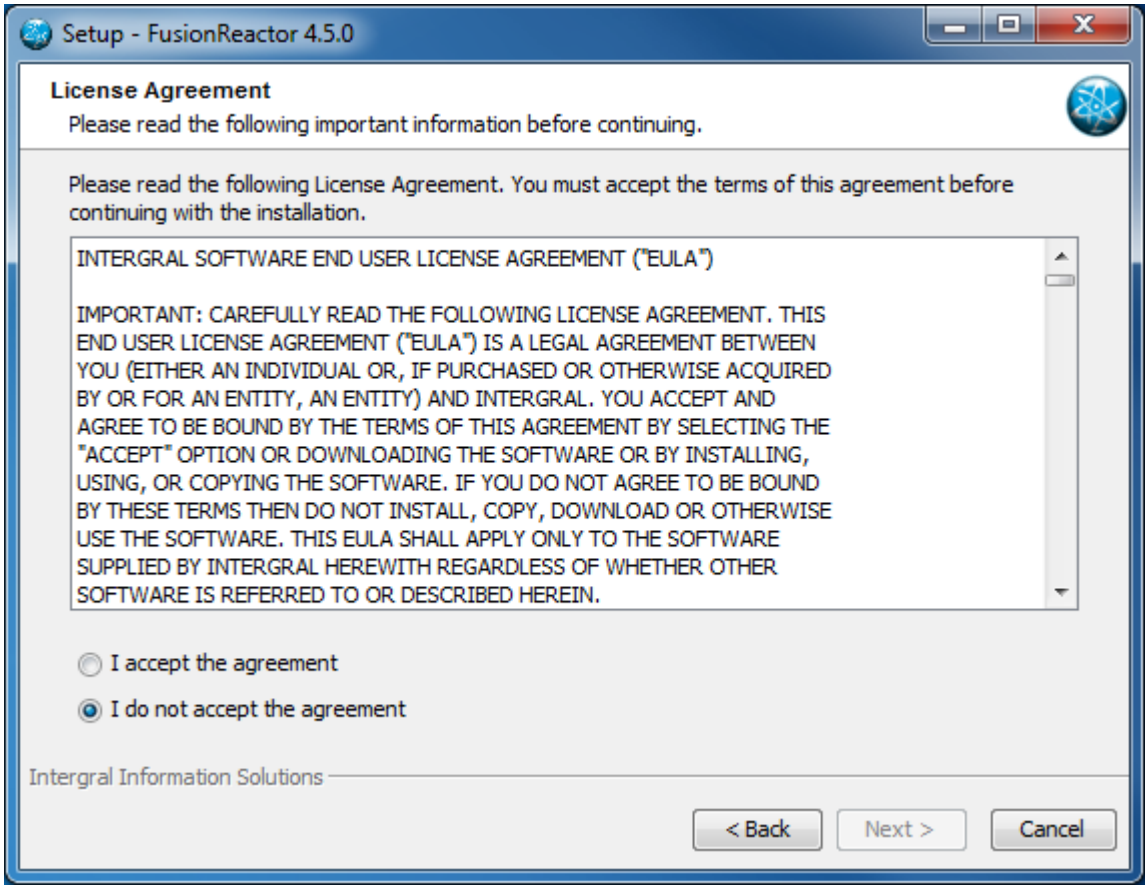


Figure 3: License Agreement

Figure 3: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.



Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the server ports (See Figure 5). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 6).

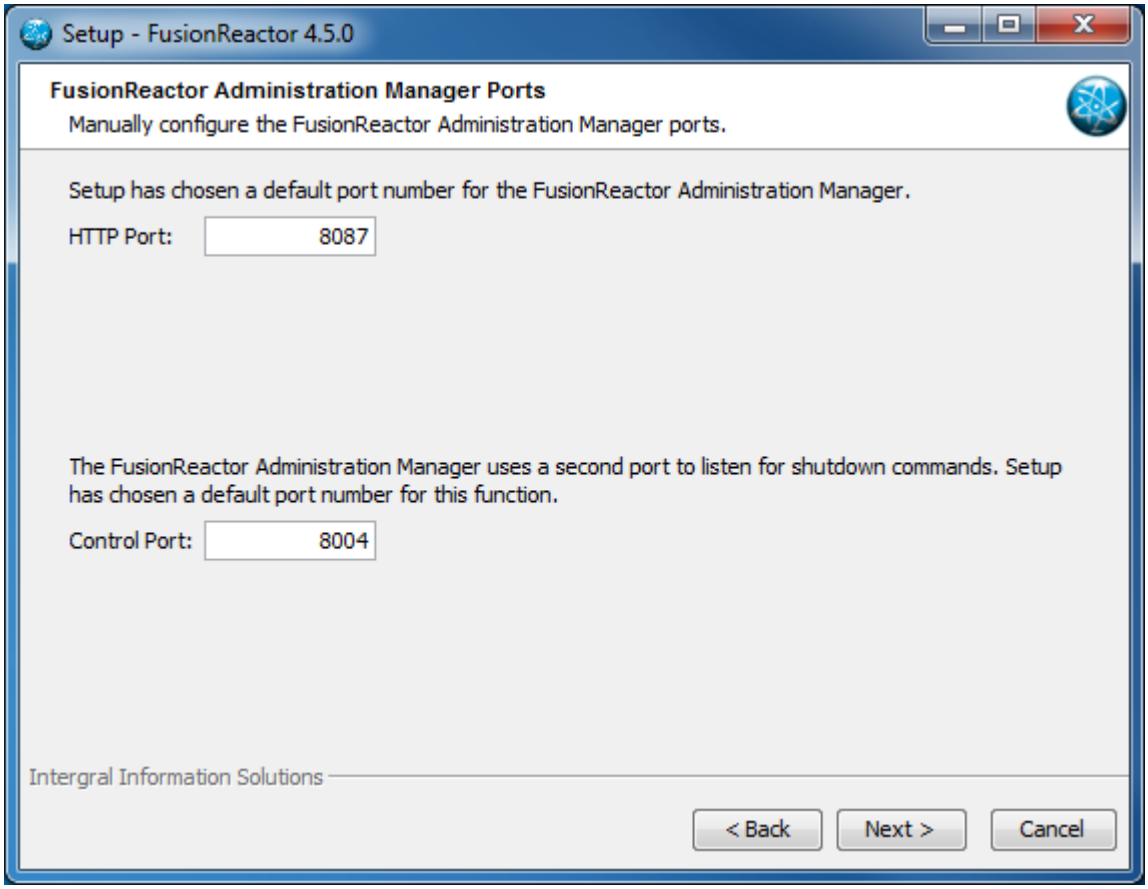


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

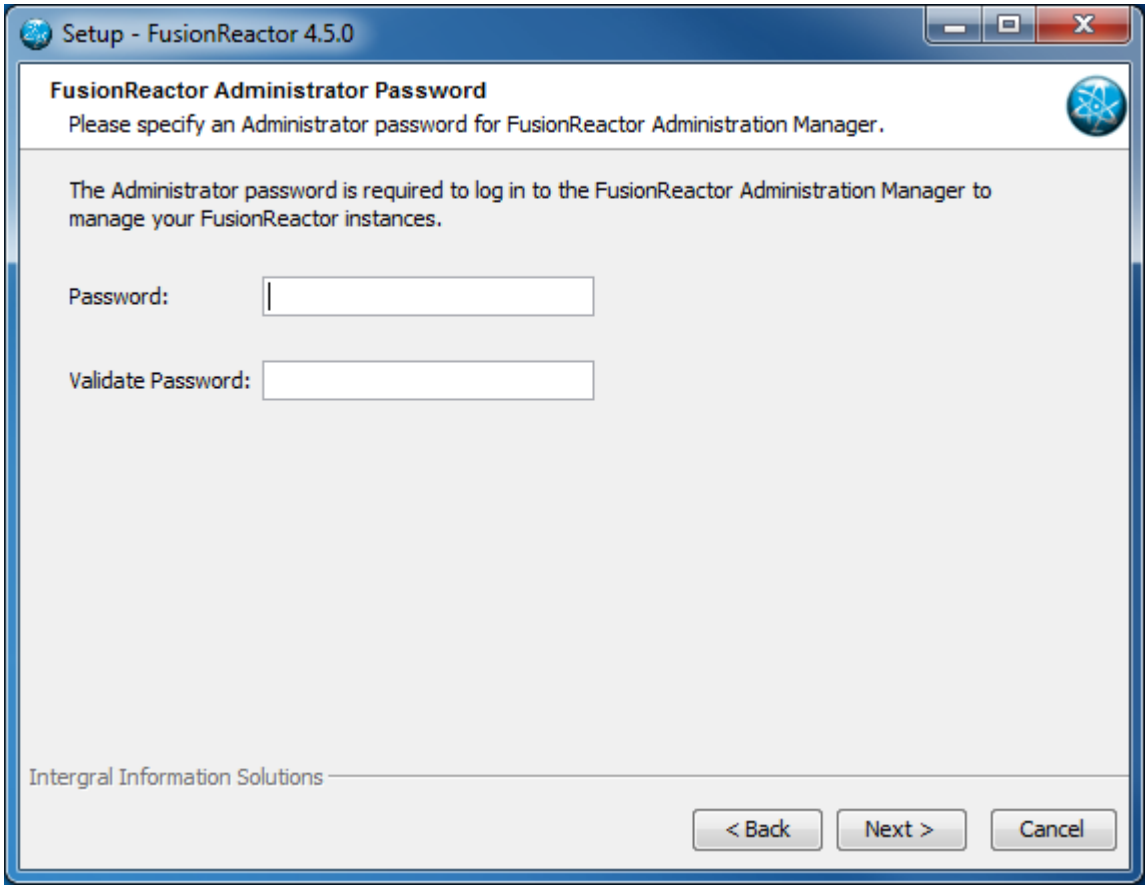


Figure 6: Administration Manager Password

Figure 6: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

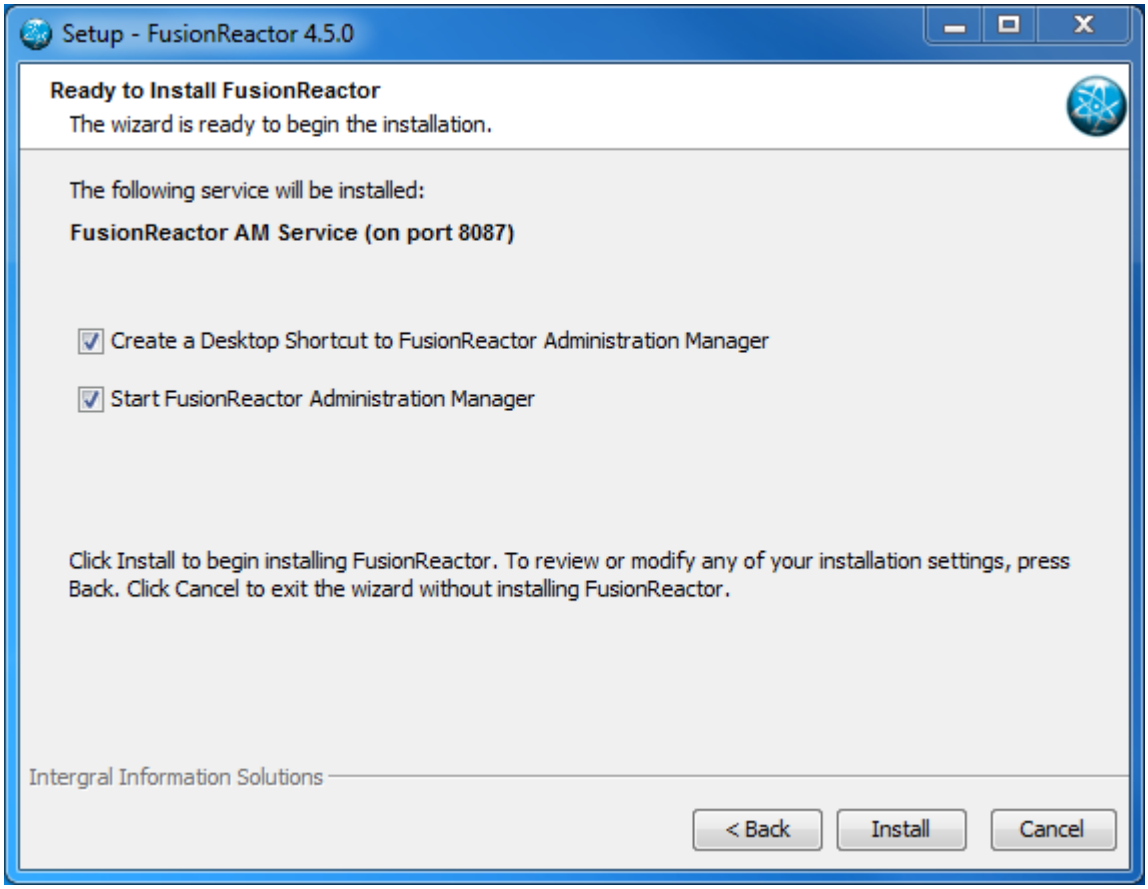


Figure 7: Ready to Install

Figure 7: The FusionReactor Setup can create a shortcut on the Desktop to FRAM, if you do not want Setup to perform this action uncheck the option *'Create Desktop Shortcut to FusionReactor Administration Manager".*

Setup can also automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin updating to FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

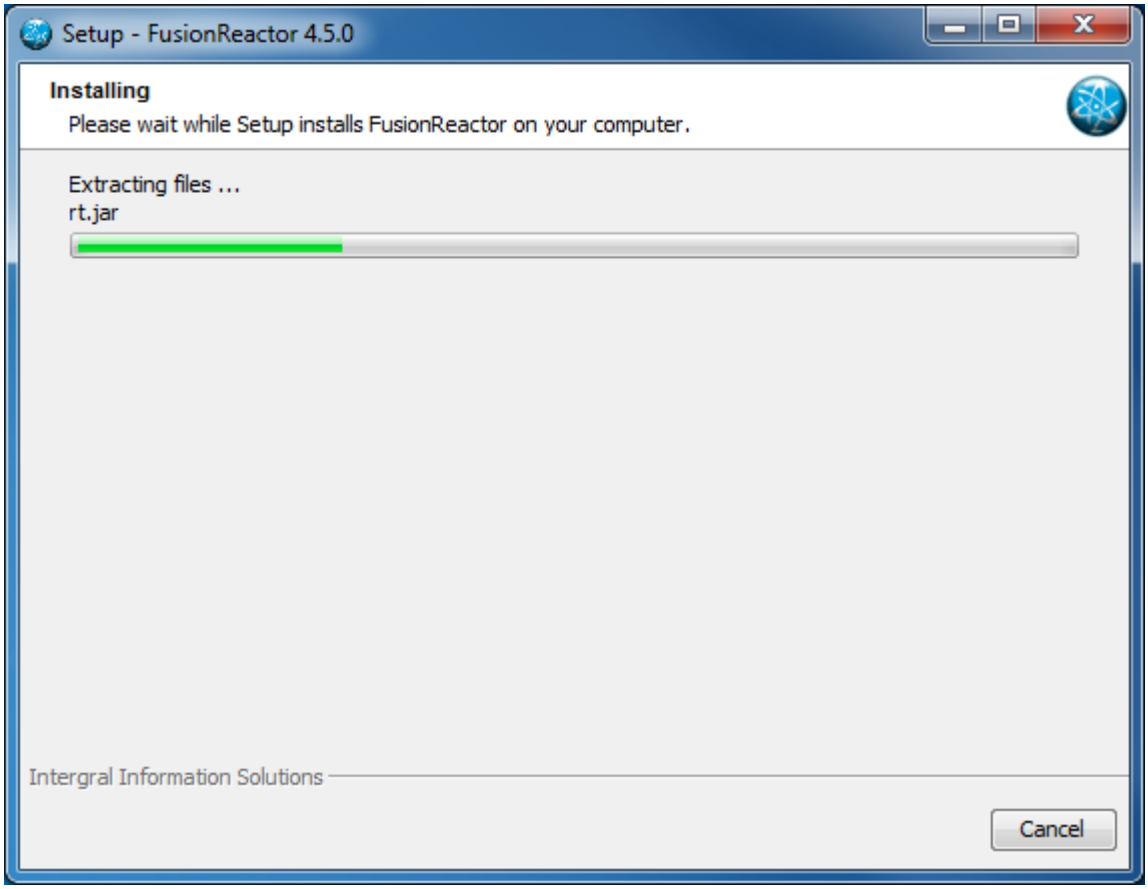


Figure 8: Installation Progress

Figure 8: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 9: Configure Instances Screen

Figure 9: After the update has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. You should now navigate to the Instance Manager located on FRAM to update your other instances of FusionReactor. Launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Update Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close'* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 7.



Figure 10: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 10: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 7. This is the last screen in the setup, clicking Finish will close the window.

Updating FusionReactor 2.0 on Windows

Upgrade Not Possible

FusionReactor 4.x.x does not support upgrades from FusionReactor 2.x.

If you currently have FusionReactor 2.x installed it is not possible to upgrade to FusionReactor 4.x.x. When running the FusionReactor 4.x.x setup you will be presented with a screen similar to that shown in Figure 1 below.

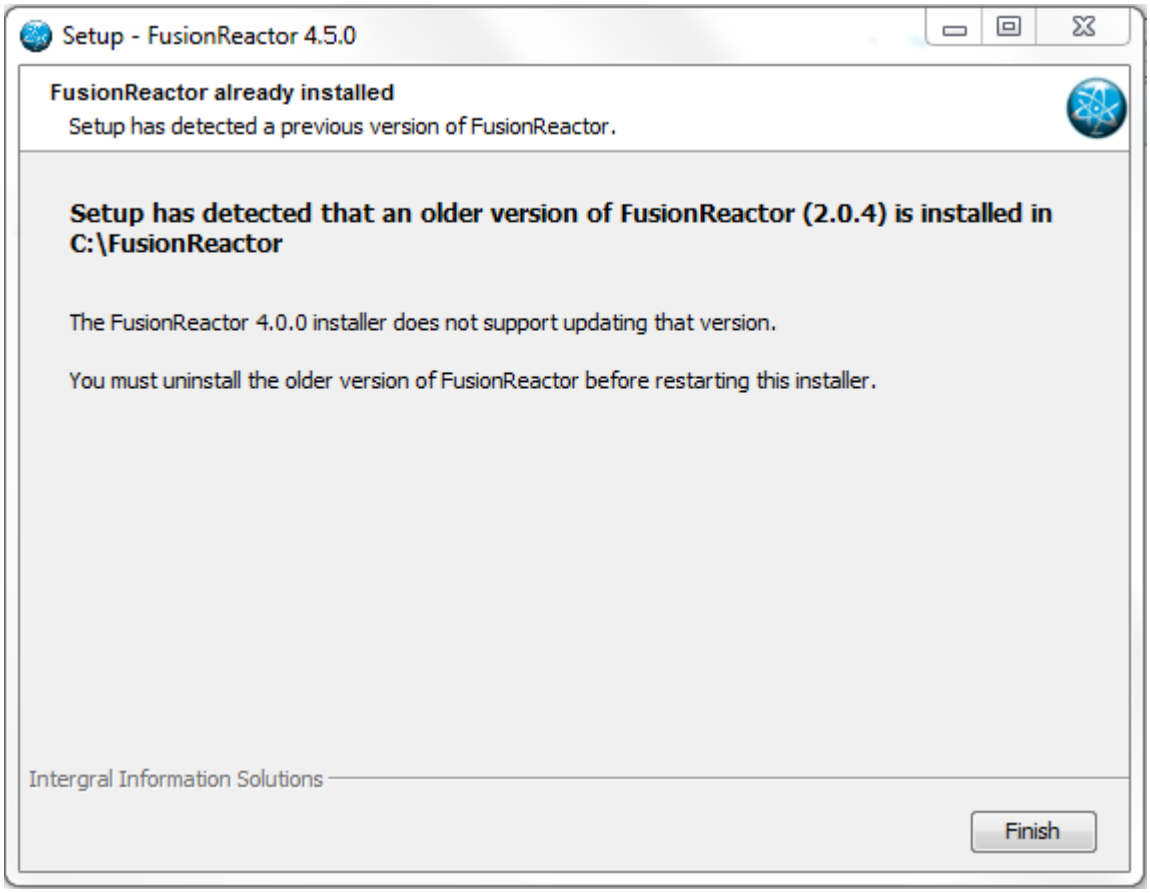


Figure 1: Update Not Supported

To install FusionReactor 4.5.x you must first uninstall FusionReactor 2.x and then run the FusionReactor 4.5.x Setup again. Please follow the guide on how to Uninstall FusionReactor 2.x on Windows

Updating FusionReactor 1.0 on Windows

> **Upgrade Not Possible**
>
> FusionReactor 4.x.x does not support upgrading from FusionReactor 1.x.

If you currently have FusionReactor 1.x installed it is not possible to upgrade to FusionReactor 4.x.x.

To install FusionReactor 4.5.x you must first uninstall FusionReactor 1.x and then run the FusionReactor 4.5.x Setup again.

Console Installation on Windows

Executables

Executable	System Architecture
FusionReactor_windows_4_5_0.exe	32-bit
FusionReactor_windows_4_5_0_x64.exe	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.

Setup

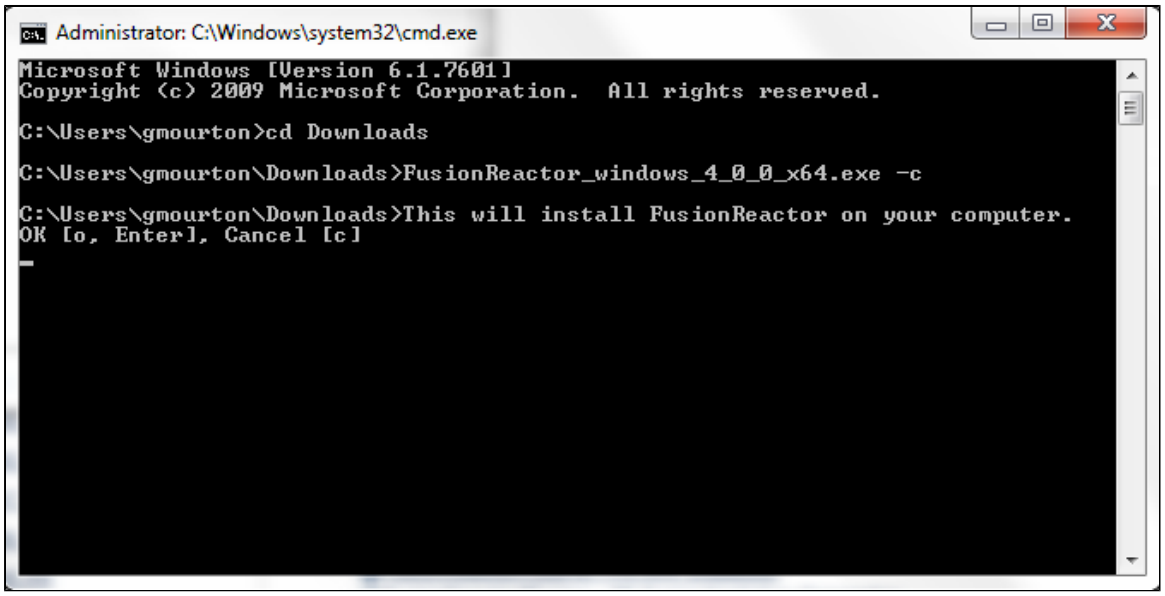


Figure 1: Start the Installation

Figure 1: To run the FusionReactor 4.5.x installation in console mode, you need to add the c switch .
For example: `FusionReactor_windows_4_5_0_x64.exe -c`

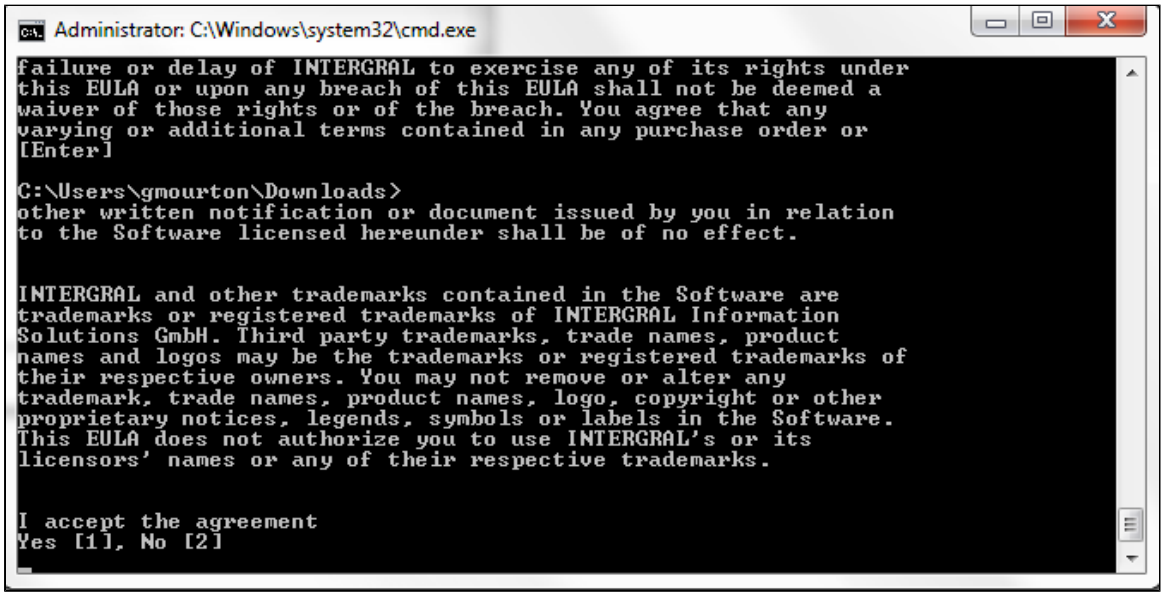


Figure 2: License Agreement

Figure 2: After starting the install you will be shown the license agreement. Please carefully review the license agreement, then enter '1' (if you accept the terms and conditions) or enter '2' to cancel the installation.

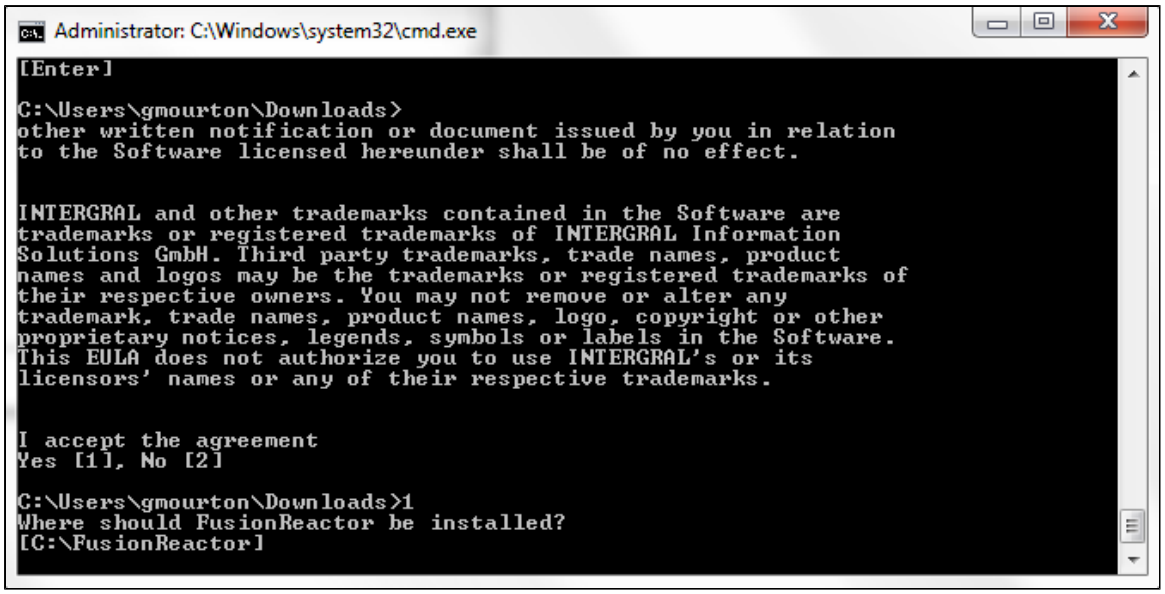


Figure 3: Installation Directory Selection

Figure 3: Now you have to select where FusionReactor is to be installed. By default (pressing enter directly), FusionReactor will be installed into `C:/FusionReactor/`.

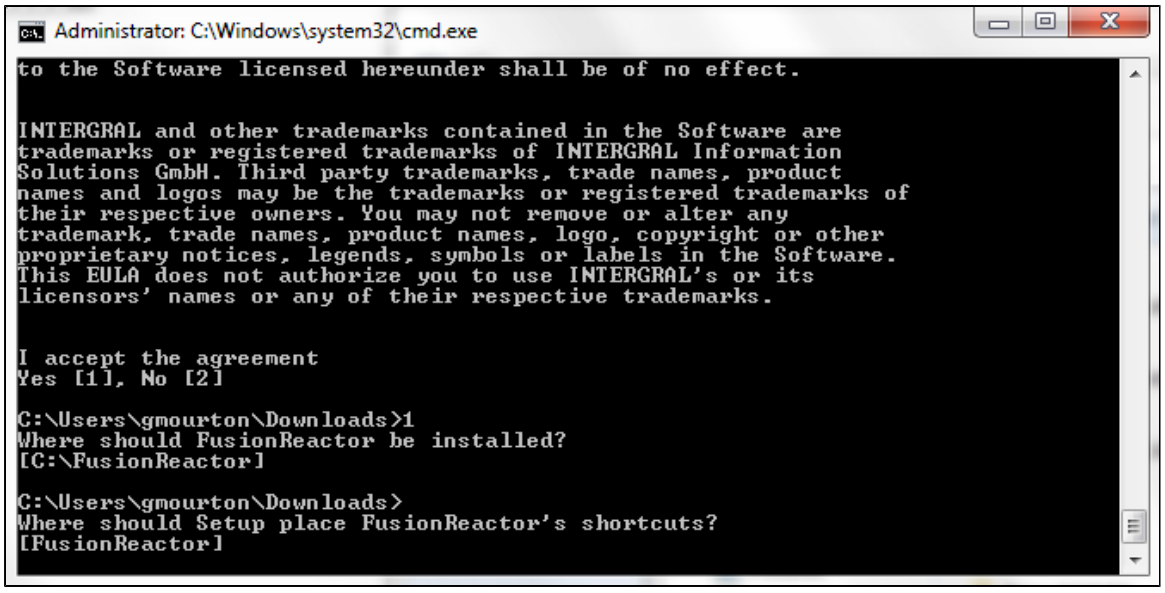


Figure 4: Shortcuts

Figure 4: You will now be asked where you want to install FusionReactors shortcuts.

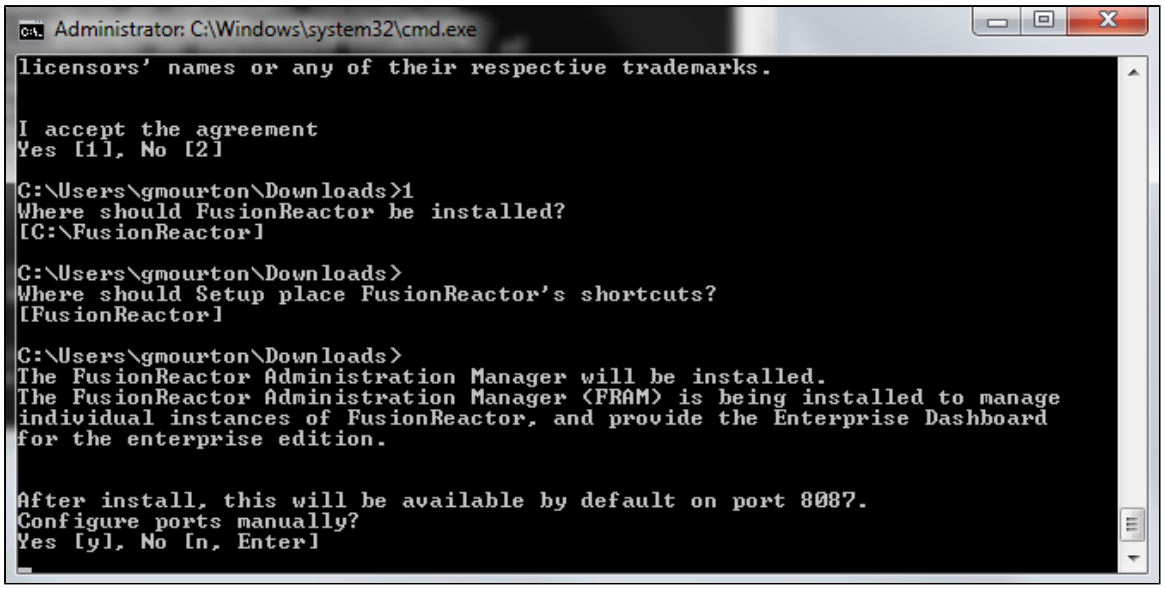


Figure 5: FusionReactor Administration Manager (FRAM)

Figure 5: You will now be asked if you want to configure the ports manually for FusionReactor Administration Manager (FRAM). Either enter 'y' to configure them (see Figure 6 onwards) or enter either 'n' or just Enter on the keyboard (see Figure 8 onwards).

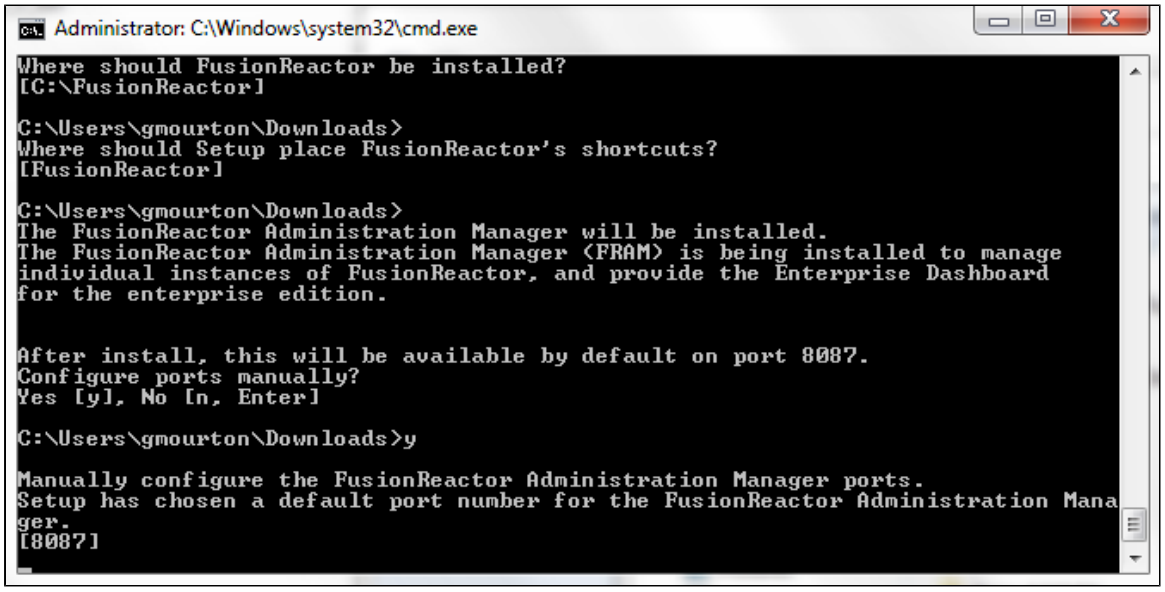


Figure 6: Manually Configuring Ports

Figure 6: Here you can configure the HTTP Port (the port you use to view FusionReactor in a web browser). It is important not to use ports that are already in use by other applications. NOTE that the default port for FRAM is 8087.

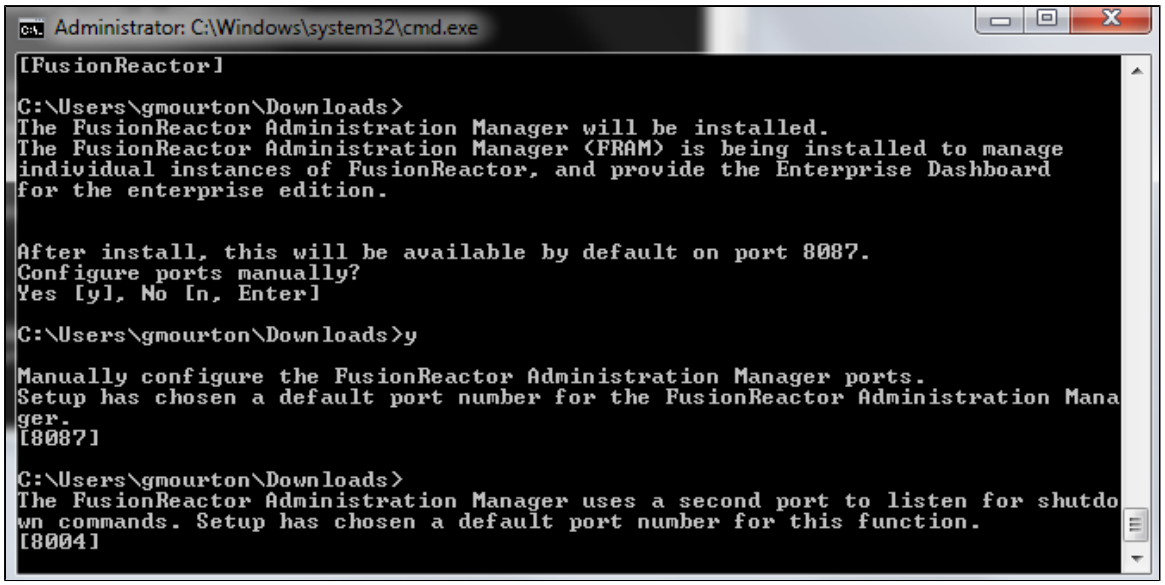


Figure 7: Manually Configuring Ports

Figure 7: You must also configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications.

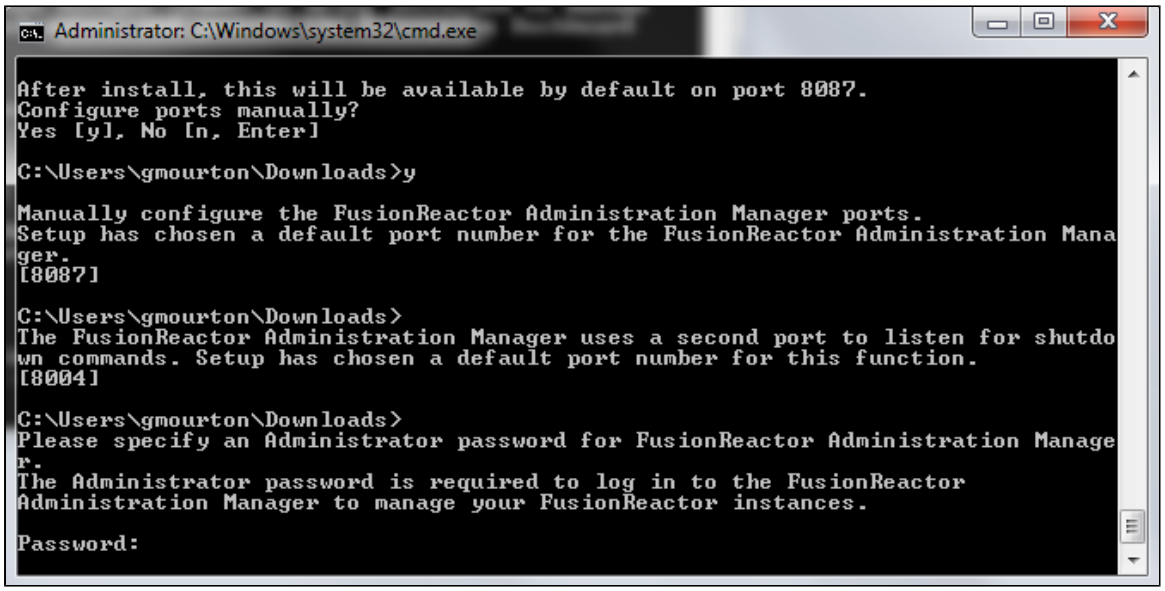


Figure 8: Administration Manager Password

Figure 8: Here you will be asked to specify a password for FRAM's Administrator account. It must be longer than 5 characters.

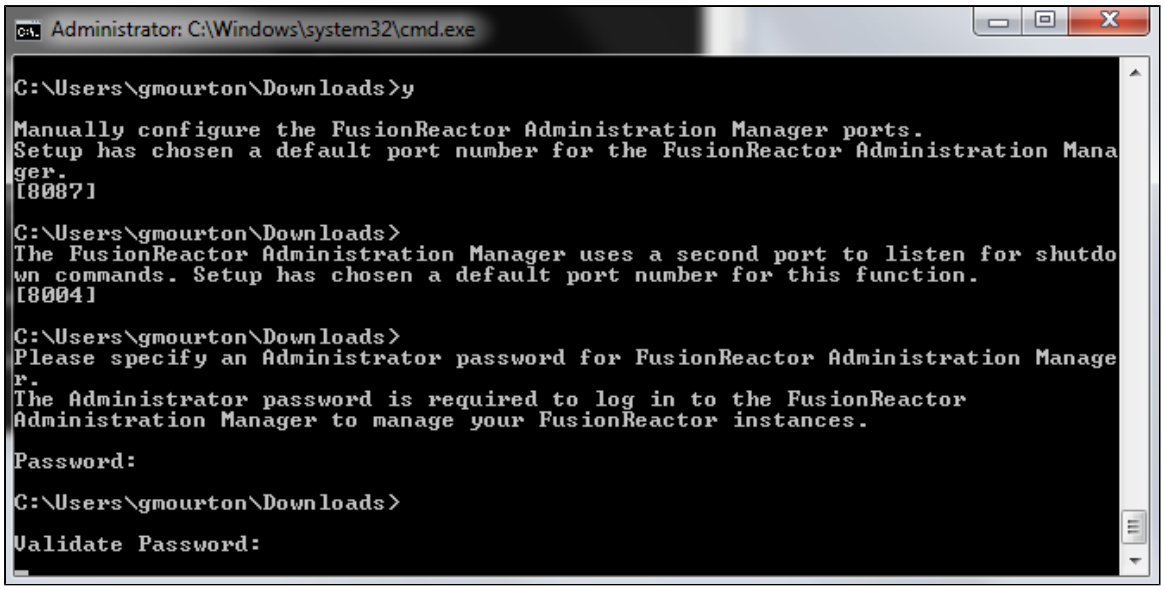


Figure 9: Administration Manager Password

Figure 9: You will be asked to validate your password.

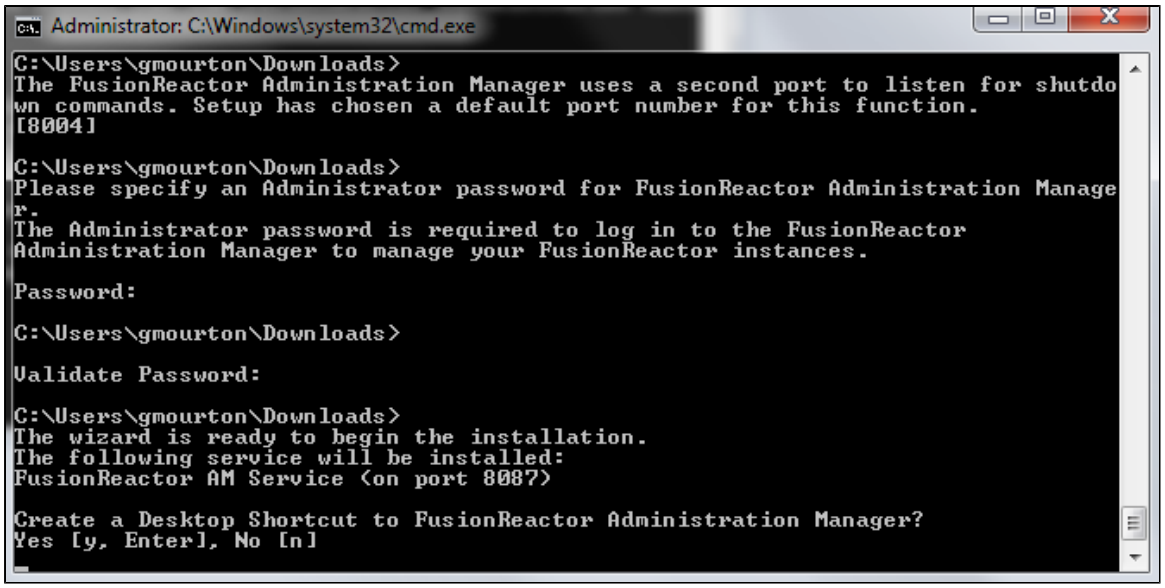


Figure 10: Create Desktop Shortcut

Figure 10: Here you will be asked whether or not you would like to create a Desktop shortcut. Pressing enter or entering 'y' will create the shortcut while pressing 'n' will not.

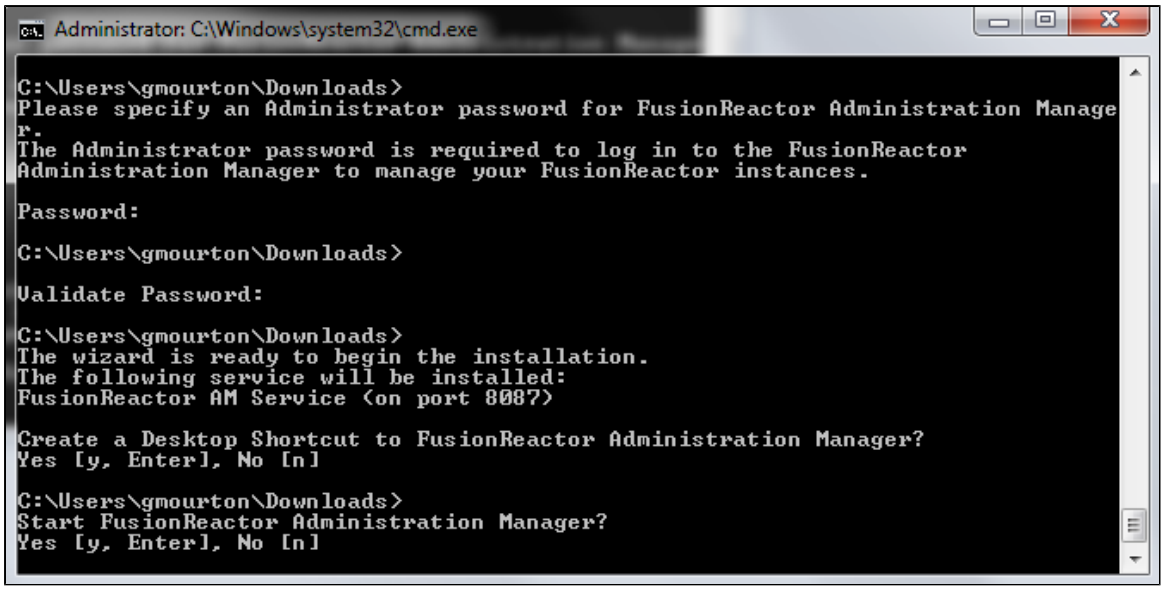


Figure 11: Ready to Install

Figure 11: Here you will be given the choice of starting up the FRAM service. Either entering 'y' or pressing enter will start the service while entering 'n' will not start the service. FRAM can be started manually after the installation has finished.

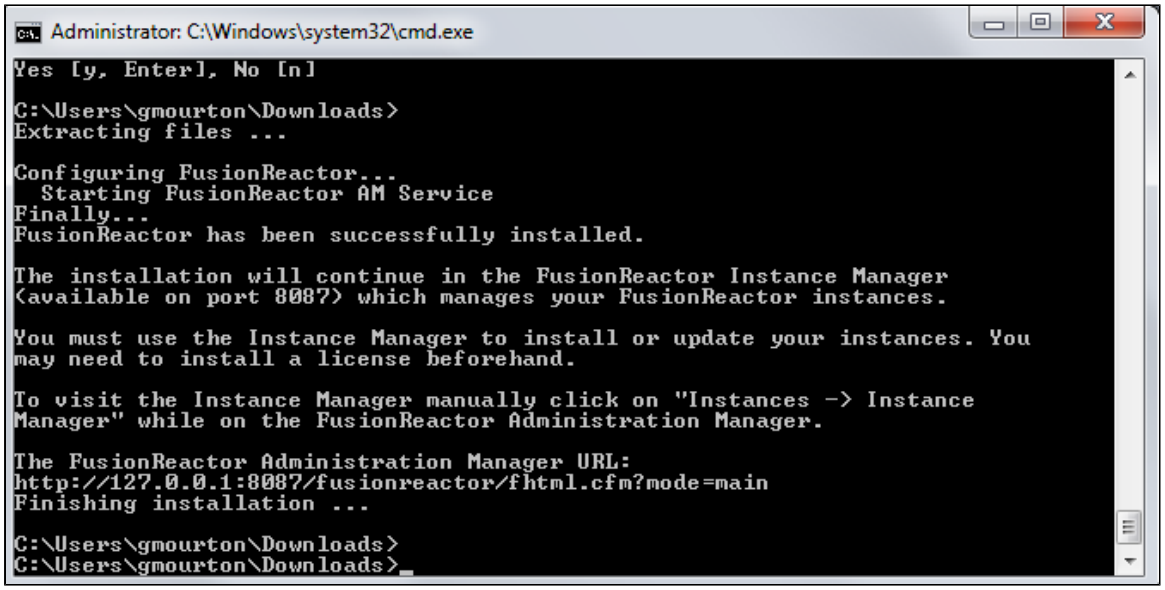


Figure 11: Finish Screen

Figure 11: You have successfully installed FusionReactor. Open the URL to go to FRAM.

Updating FusionReactor 4.0.x on Windows

Executables

Executable	System Architecture
FusionReactor_windows_4_5_0.exe	32-bit
FusionReactor_windows_4_5_0_x64.exe	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

You can run the FusionReactor 4.5.x Setup to update an already installed version of FusionReactor 4.0.x

To start the update, login as user *Administrator* and run the appropriate executable as described in the above section **Executables.**



Figure 1: Welcome Screen

Figure 1: The first step in the installation wizard is the welcome screen as shown in Figure 1. Clicking the Next button will proceed with the update.

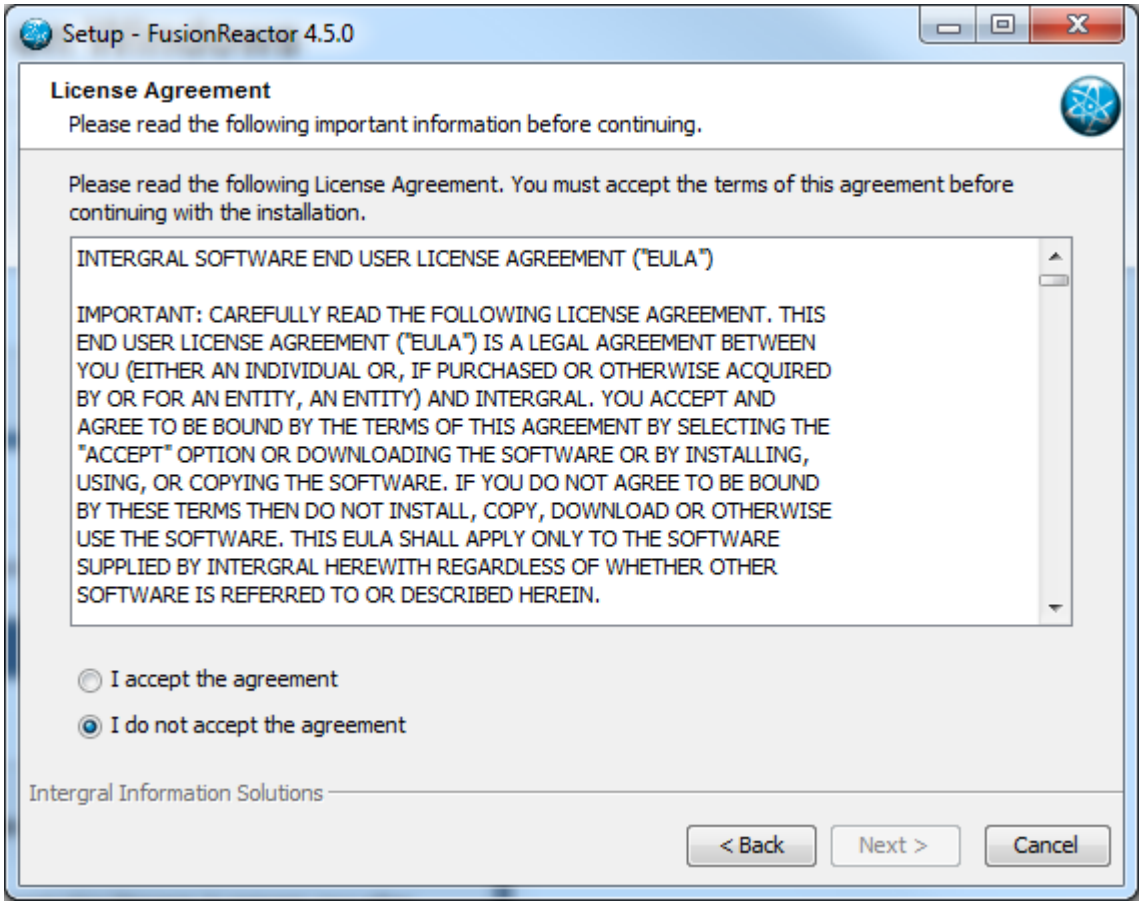


Figure 2: License Agreement

Figure 2: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.

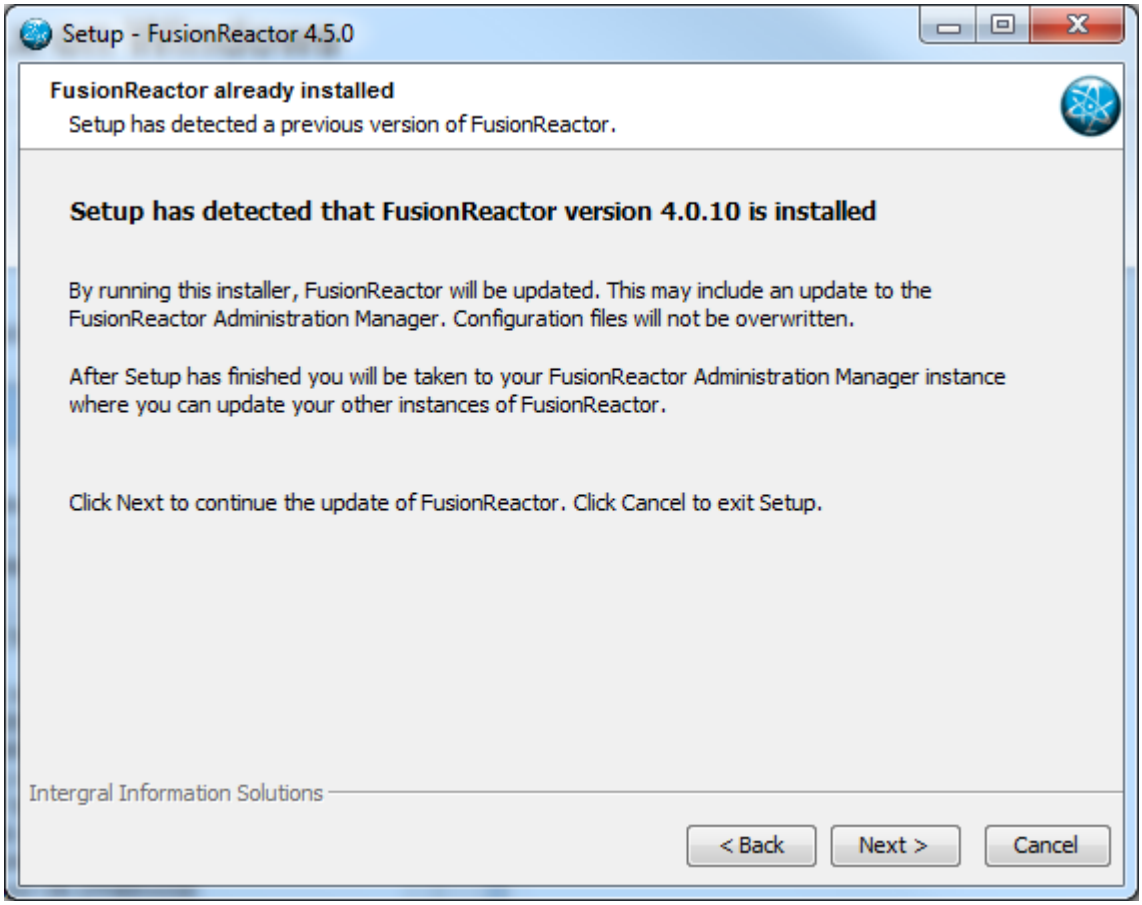


Figure 3: Update the Current Version

Figure 3: If there is already a version of FusionReactor 4.0.x on your computer that has been installed with the FusionReactor 4.0.x Setup, you will see a screen similar to the one shown. Carefully read the information and click Next to proceed with the update. Click Cancel to exit the installation wizard.

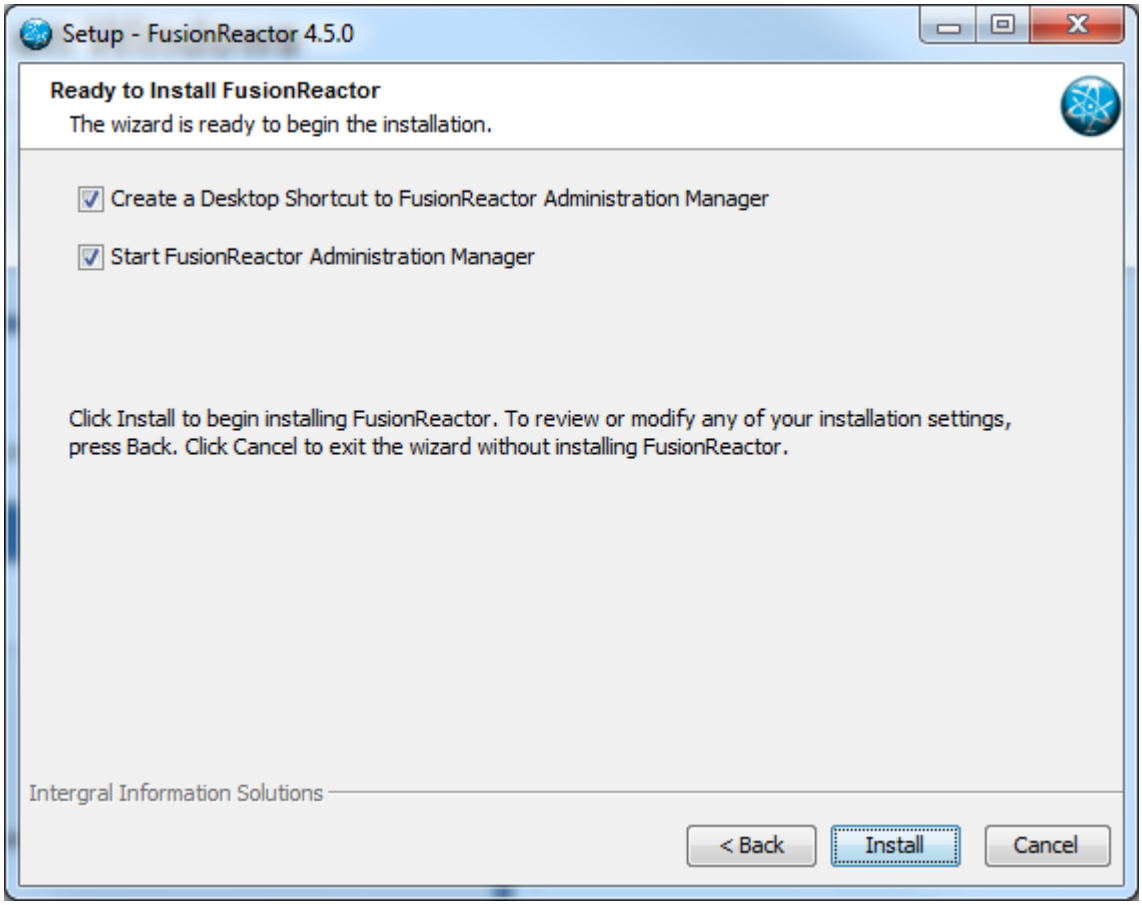


Figure 4: Ready to Install

The FusionReactor Setup can create a shortcut on the Desktop to FRAM, if you do not want Setup to perform this action uncheck the option *'Create Desktop Shortcut to FusionReactor Administration Manager".*

Setup can also automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin updating to FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

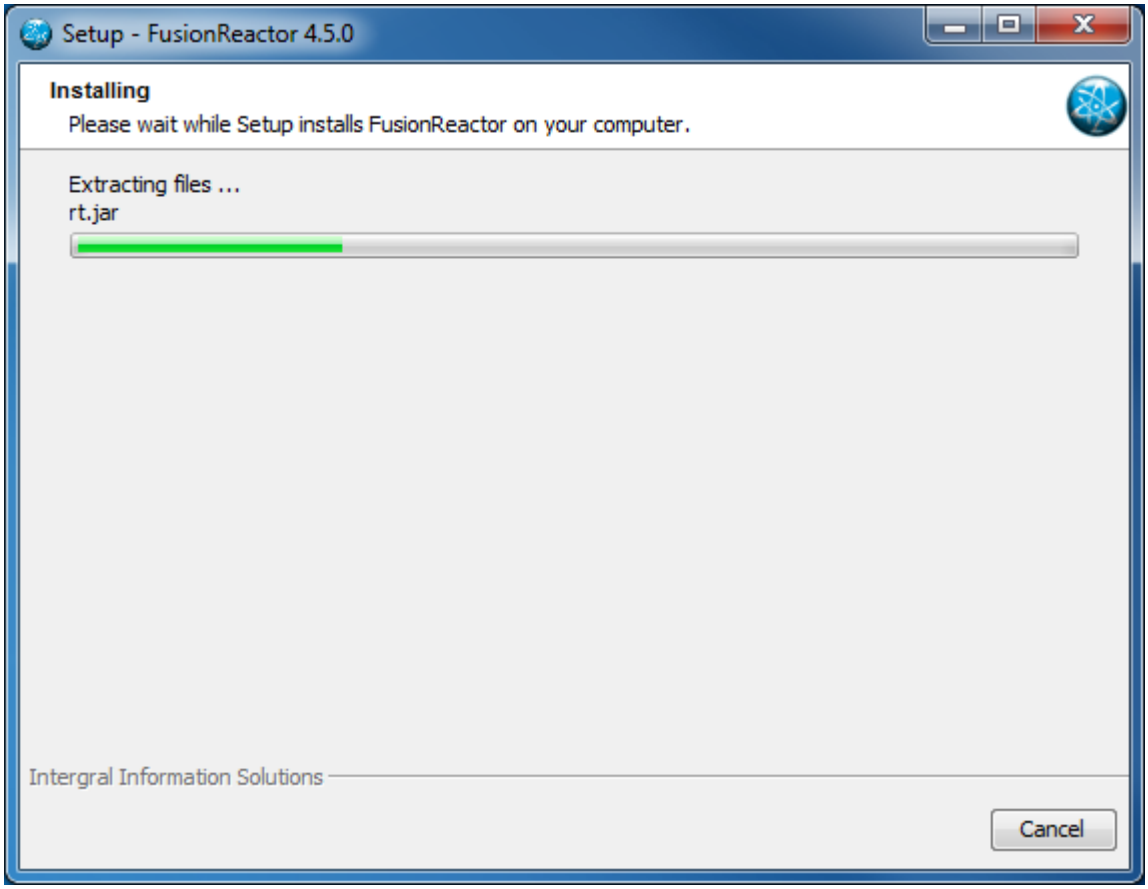


Figure 5: Installation Progress

Figure 5: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.

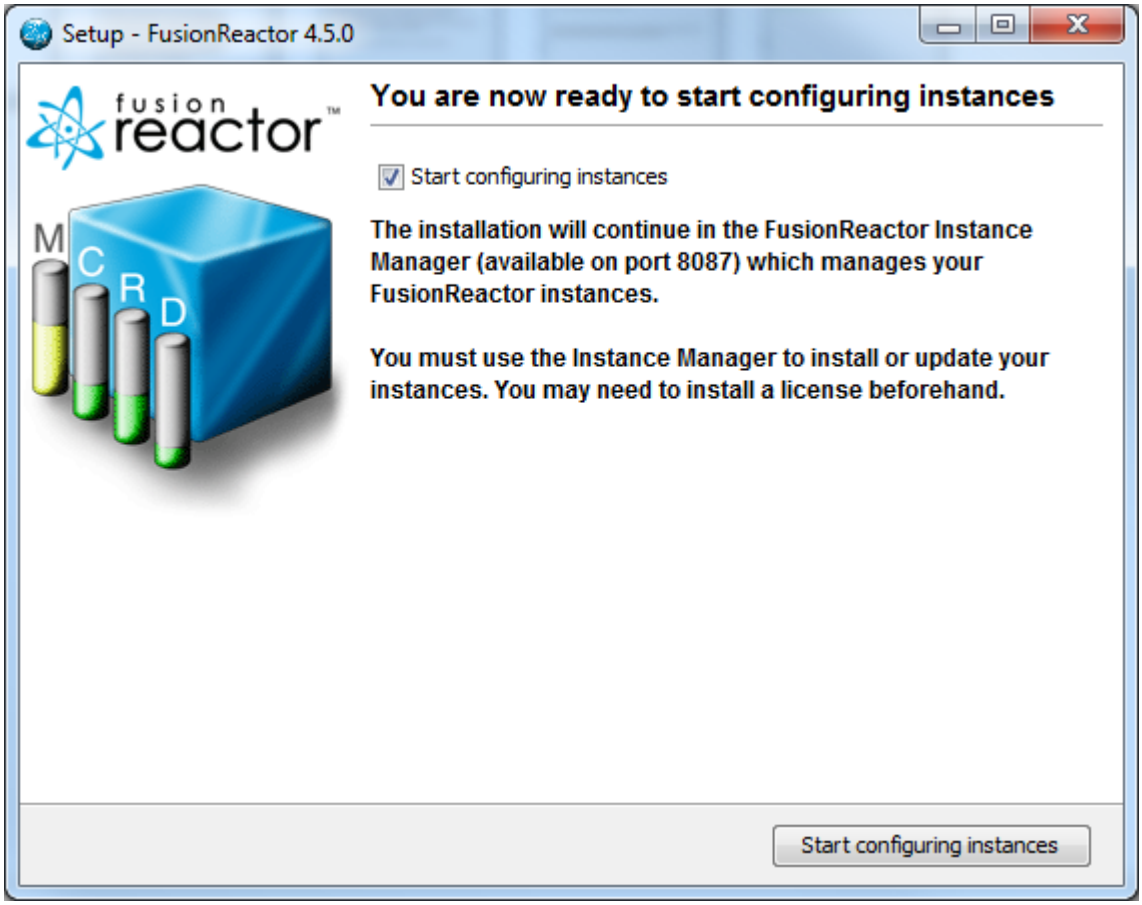


Figure 6: Configure Instances Screen

Figure 6: After the update has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. You should now navigate to the Instance Manager located on FRAM to update your other instances of FusionReactor. Launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Update Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close'* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 4.



Figure 7 Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 7: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 4. This is the last screen in the setup, clicking Finish will close the window.

Uninstalling FusionReactor on Windows

To completely remove FusionReactor from your system, you can either use the FusionReactor Uninstaller Start Menu entry or you can use the Windows Control Panel to open the Add or Remove Programs application shown below. The Uninstaller will completely remove the FusionReactor directory including all of the generated data – **p lease take care that you save any files you want to keep before running the uninstaller.**

> The uninstaller will also remove any installed FusionReactor instances, during this process you will be asked to restart any running servers. Please note that not stopping the servers means that FusionReactor and all of its components will not be completely removed. This may cause issues on these servers.

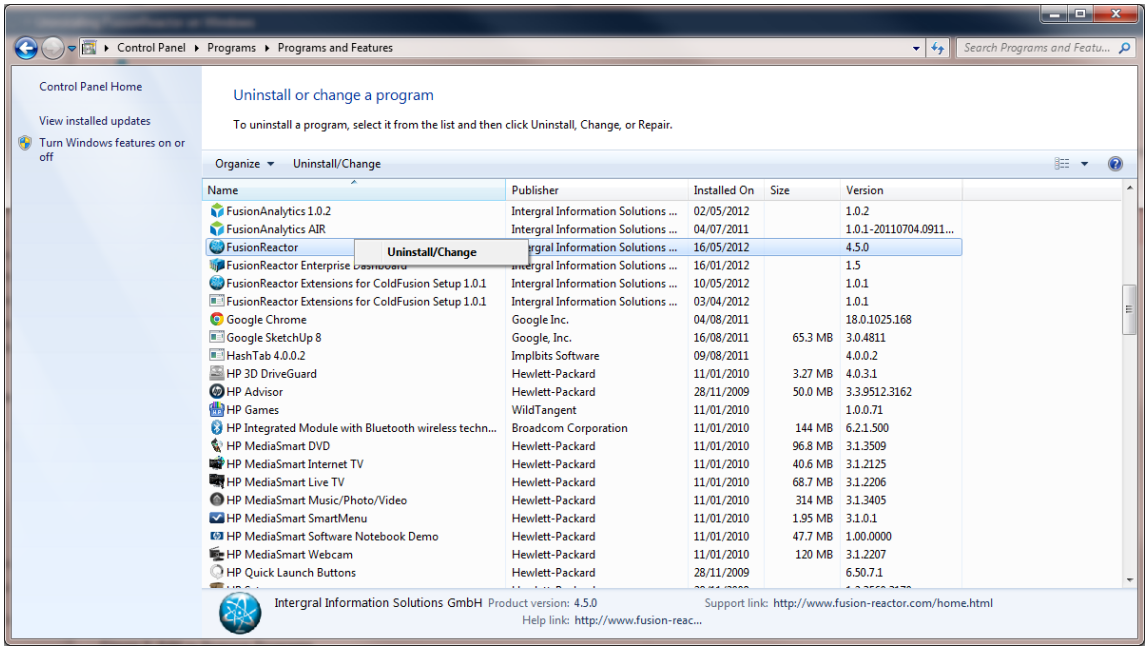


Figure 1: Add or Remove Programs

Figure 1: Navigate through the Add or Remove Programs application and find FusionReactor 4.0.x. If you are sure you want to uninstall FusionReactor, right click the FusionReactor entry you found and right click uninstall/change.

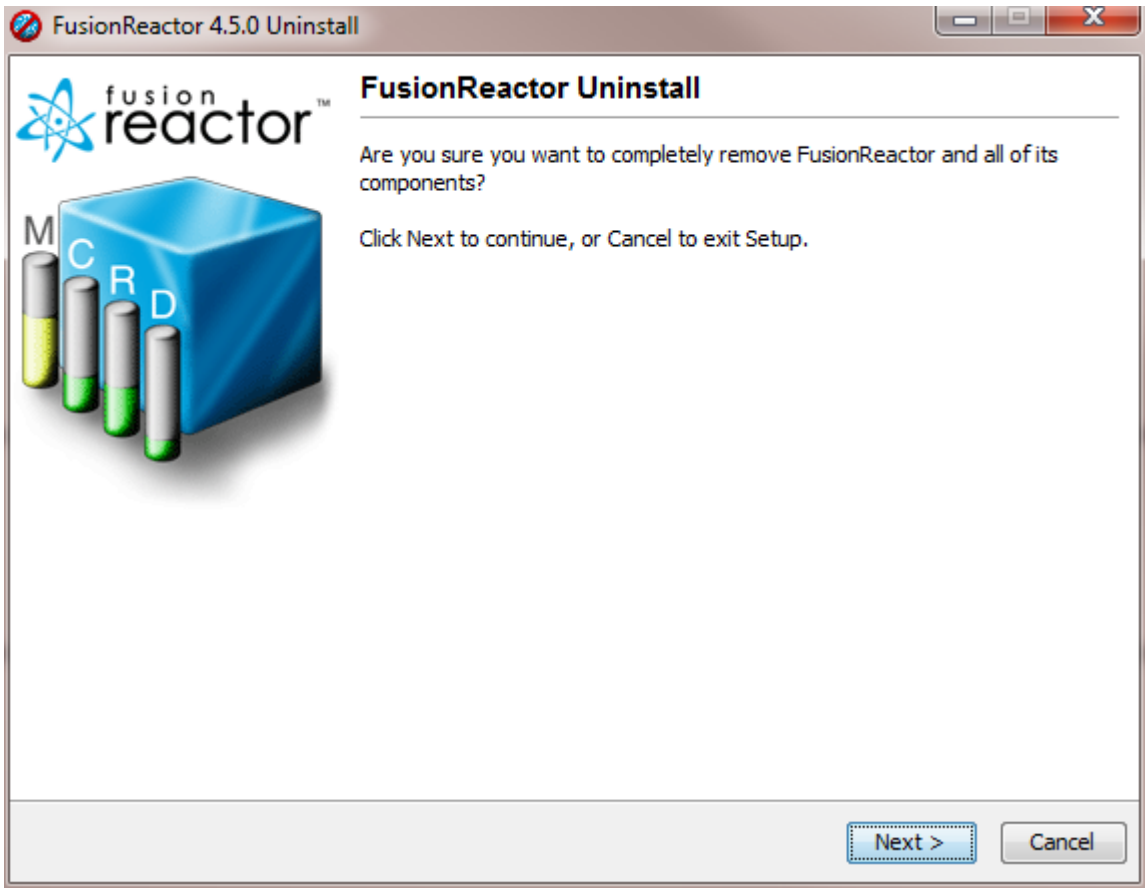


Figure 2: FusionReactor Uninstall

Figure 2: Once the Uninstaller has been started, read the text that asks you if you want to remove FusionReactor and all it's components. Click Next to continue on with the uninstall or click Cancel to exit the

uninstaller wizard.

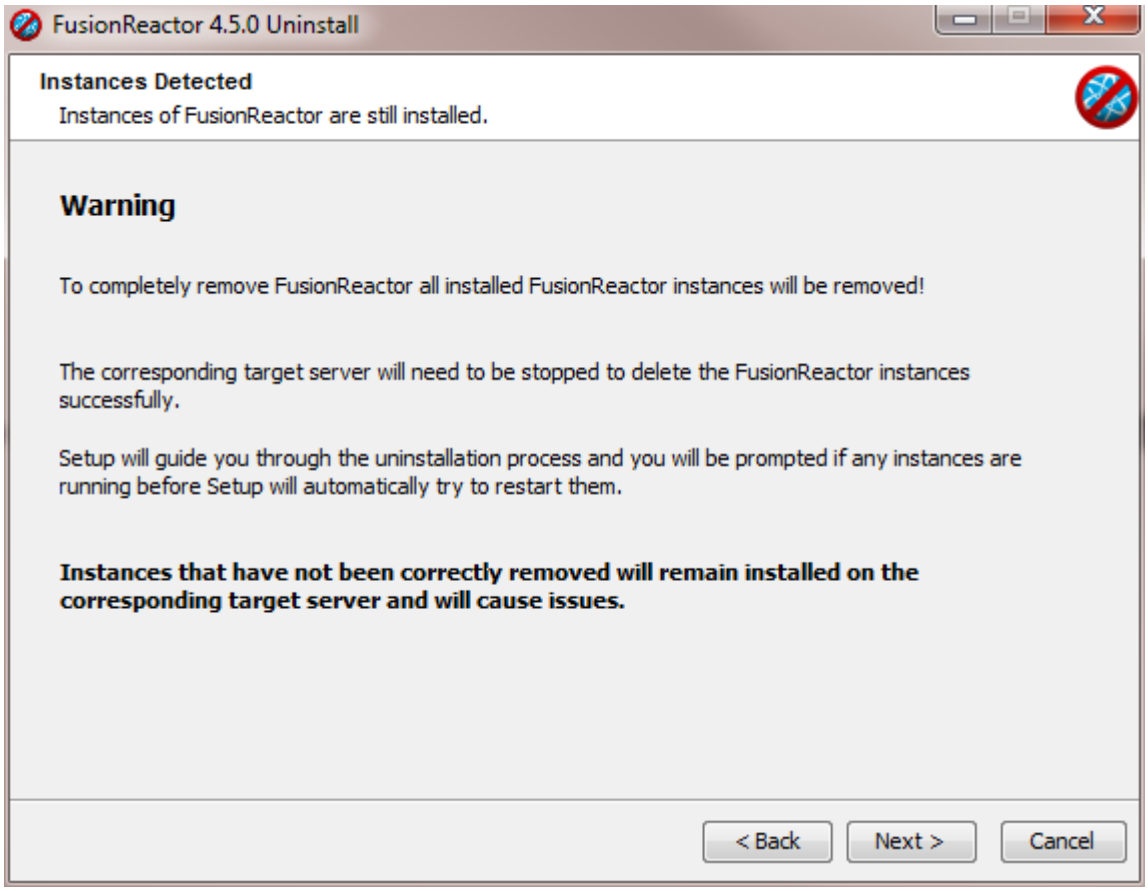


Figure 3: Uninstall FusionReactor Instances Warning

Figure 3: If you have other FusionReactor instances installed a warning message will be displayed to highlight the effects of removing FusionReactor. Carefully read through this. Once you have read it and you have understood it, please **click Next to uninstall all of the installed FusionReactor instances as well as FusionReactor.** Otherwise, click Cancel to stop the uninstall.

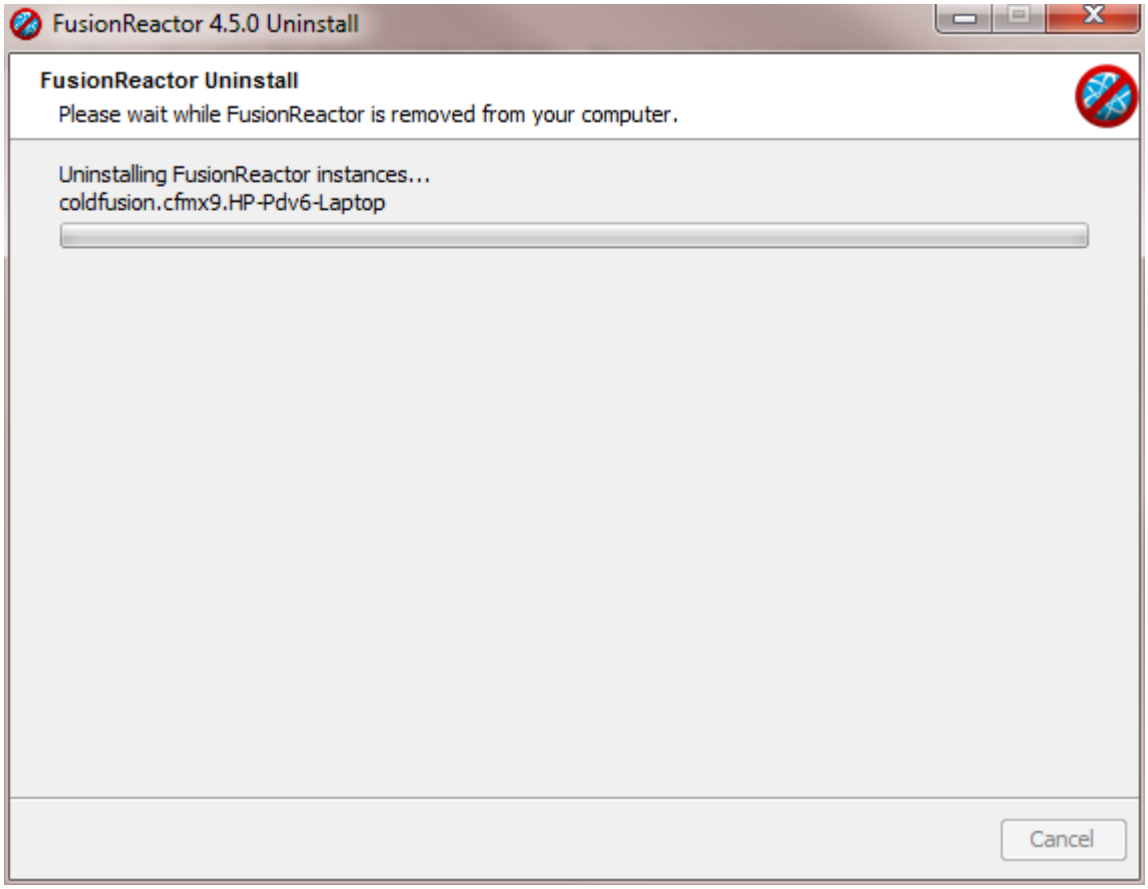


Figure 4: Uninstall Progress Bar

Figure 4: Once you have clicked Next, the FusionReactor uninstaller will begin to uninstall FusionReactor and all of the installed FusionReactor instances.

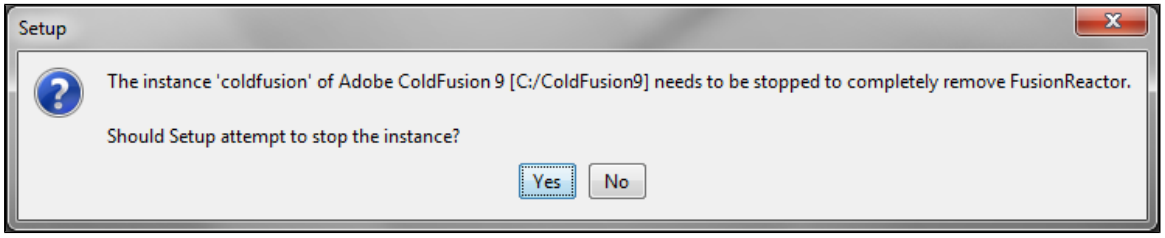


Figure 5: Stopping a Server Instance

Figure 5: If there is a FusionReactor instance that is running, the server container for this needs to be stopped in order to successfully remove the FusionReactor Instance.

Press Yes if Setup should automatically attempt to stop the server instance, No if you wish to leave the server running.

As stated at the top of this page, not stopping the server container means that FusionReactor and all of its components will not be completely removed from the target server. This may cause issues on the corresponding server.

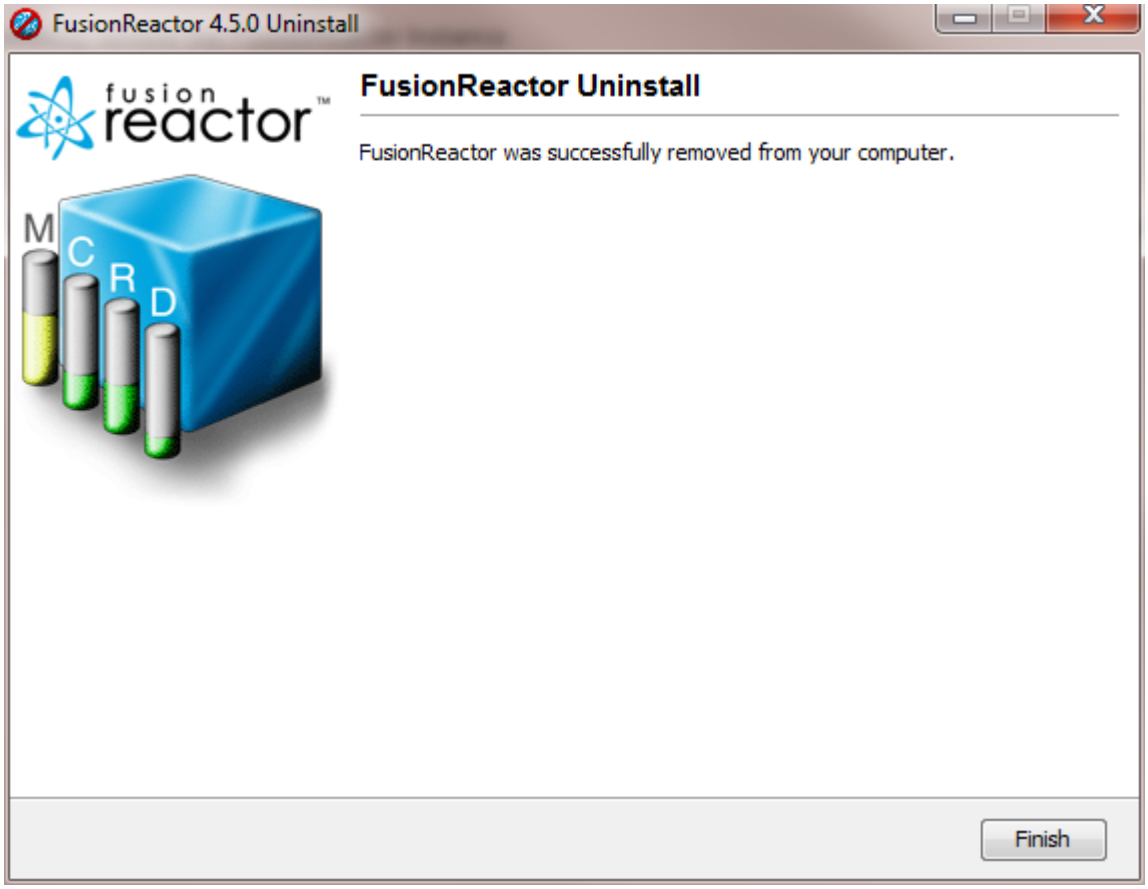


Figure 6: FusionReactor Uninstall Complete

Figure 6: When FusionReactor has finished uninstalling you will be presented with the Finish screen shown in Figure 6. Click on Finish to exit.

Installing FusionReactor on Linux

[Installing FusionReactor 4.5.x on Linux]

[Updating FusionReactor 3.x on Linux]

[Updating FusionReactor 2.0 on Linux]

[Updating FusionReactor 1.0 on Linux]

[Console Installation on Linux]

[Updating FusionReactor 4.0.x on Linux]

Installing FusionReactor 4.5.x on Linux

Executables

Executable	System Architecture
FusionReactor_linux_4_5_0.sh	32-bit
FusionReactor_linux_4_5_0_x64.sh	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

To start the installation, login as user *root* and run the appropriate executable as described in the above section **Executables.**
If you can not use an X-Windows session, you can install FusionReactor through the console installation. Please follow the guide that covers Console Installation on Linux



Figure 1: Welcome screen

Figure 1: The first step in the installation wizard is the welcome screen.

Clicking the Next button will proceed with the installation and open the License Agreement as shown in Figure 2. Clicking Cancel will exit the Setup.

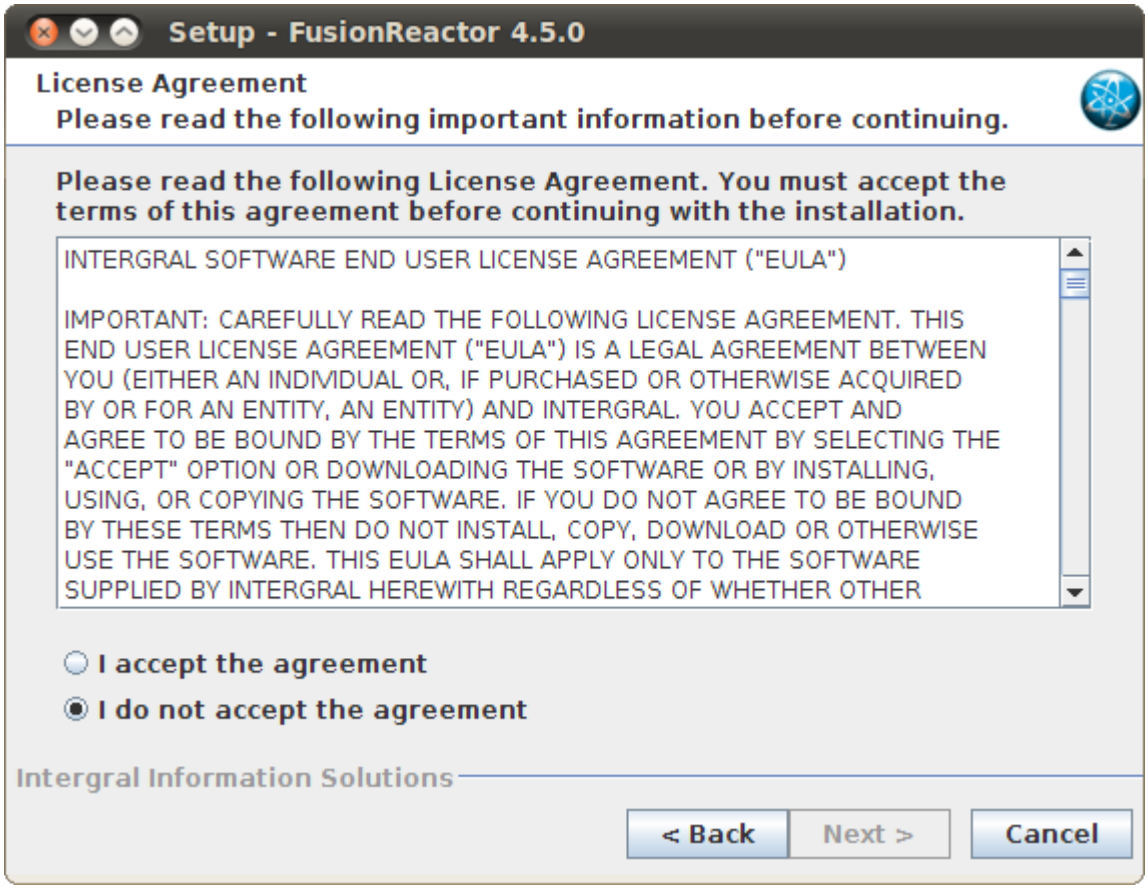


Figure 2: License Agreement

Figure 2: Please carefully review the license agreement, then choose the *'I accept the agreement'* option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.

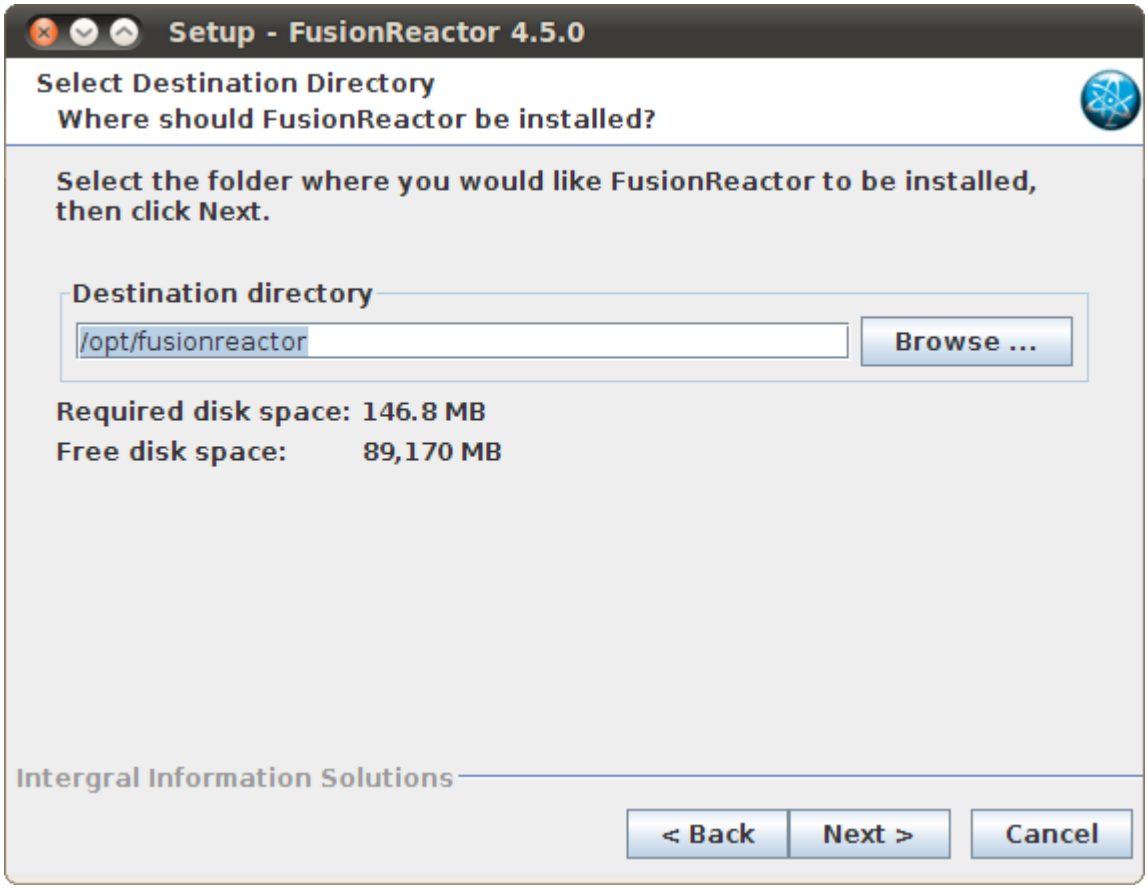


Figure 3: Installation Directory Selection

Figure 3: Now you have to select the directory in which FusionReactor should be installed. By default FusionReactor will be installed into `/opt/FusionReactor/`

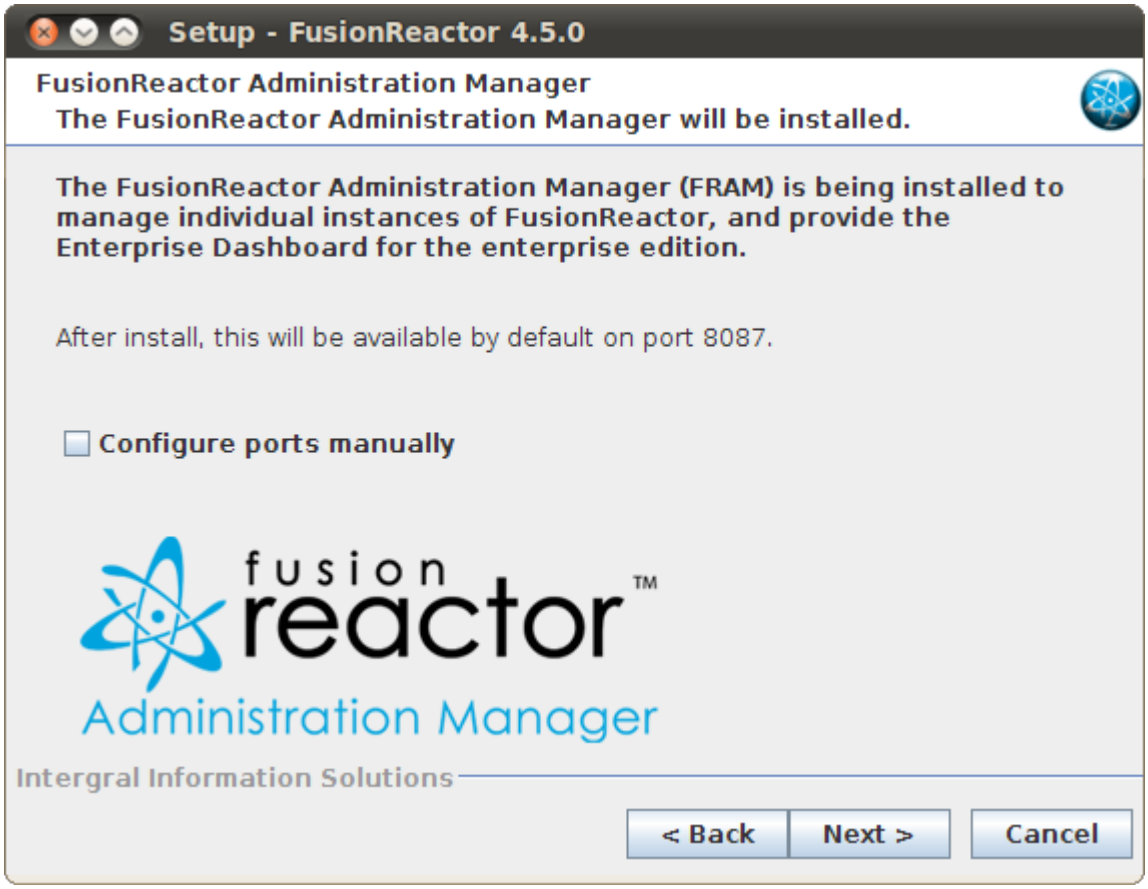


Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the sever ports (See Figure 5). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 6).

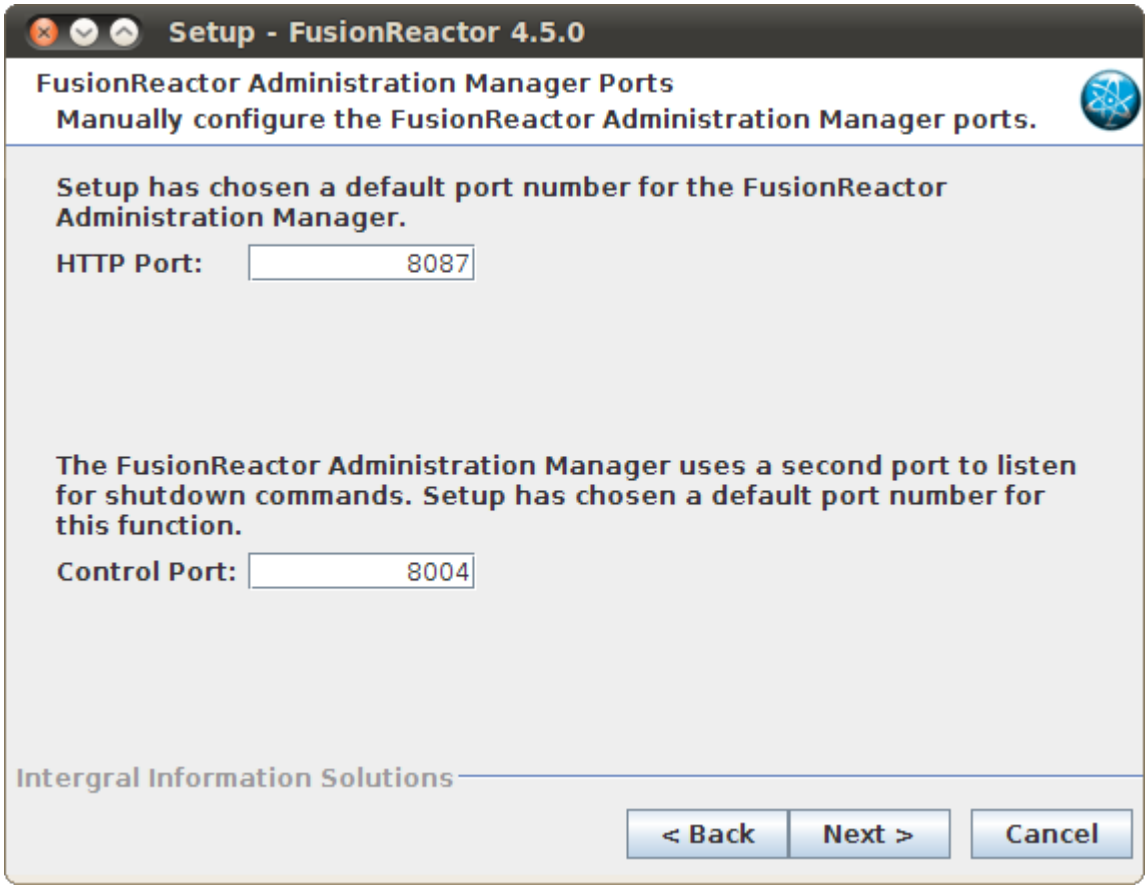


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

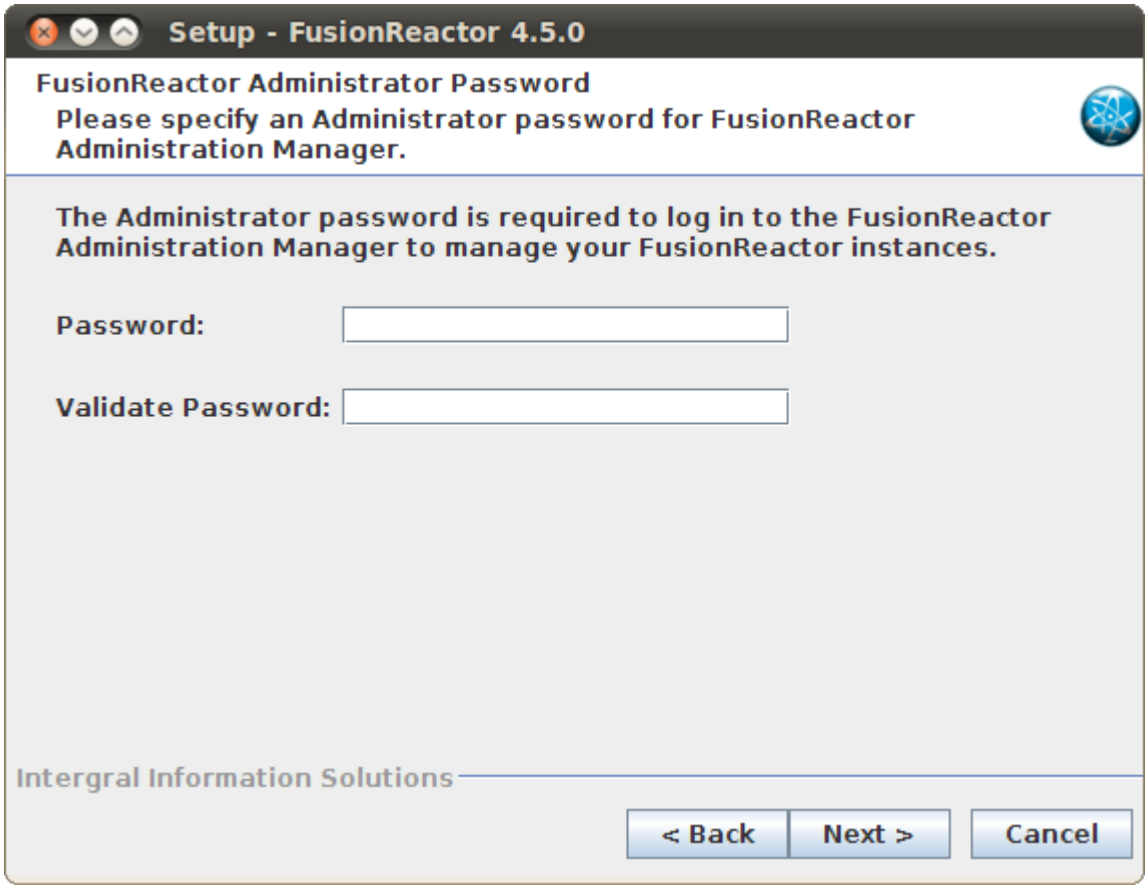


Figure 6: Administration Manager Password

Figure 6: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

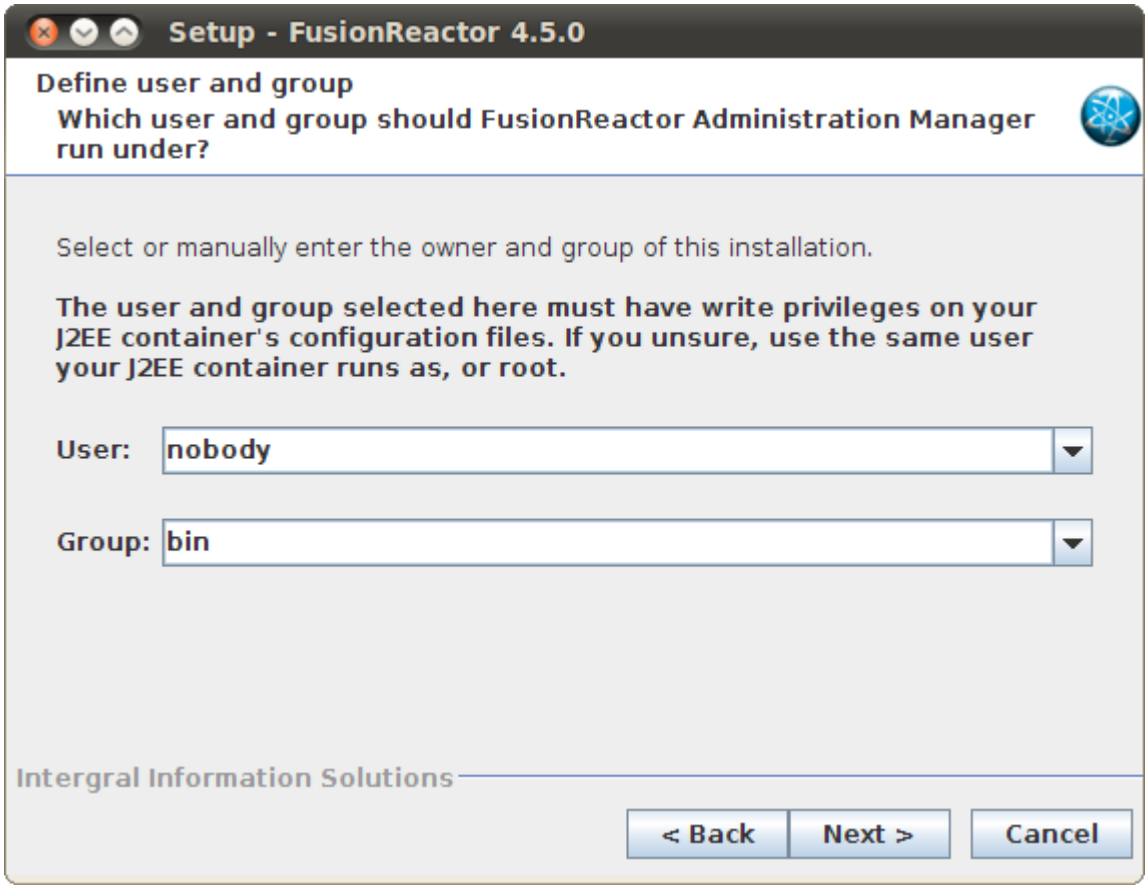


Figure 7: Choose a User and a Group

Figure 7: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM). Only that user or someone with the same or higher privileges will be allowed to run FRAM.

Please Note

FusionReactor Administration Manager (FRAM) must be installed as a user who has a real shell and can log in to the system. By default, the installer will pick **nobody/root** for this purpose, but can't presently check that this user can actually log in. If your **nobody** user does not have a shell, please change their shell to **/bin/sh** (or your preferred shell) or select another user.

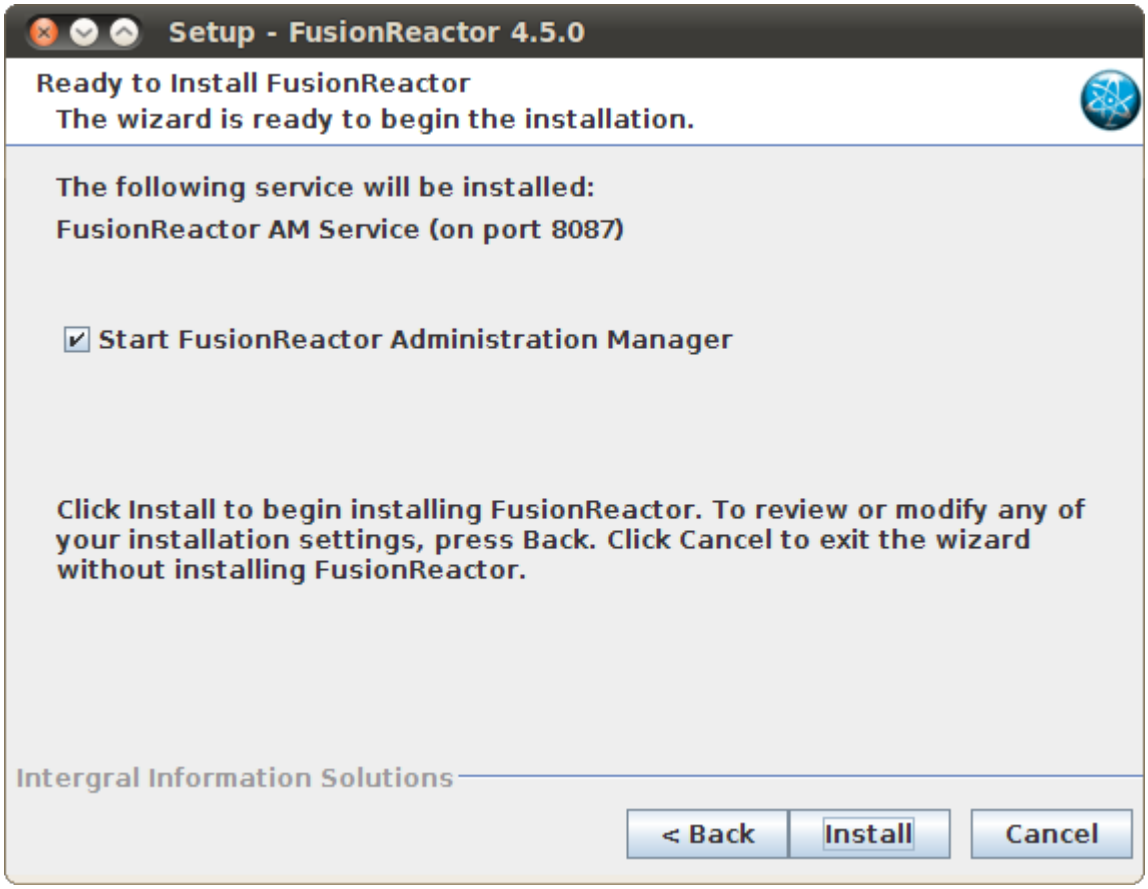


Figure 8: Ready to Install

Figure 8: Setup can automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager'*.

Click Install to begin installing FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

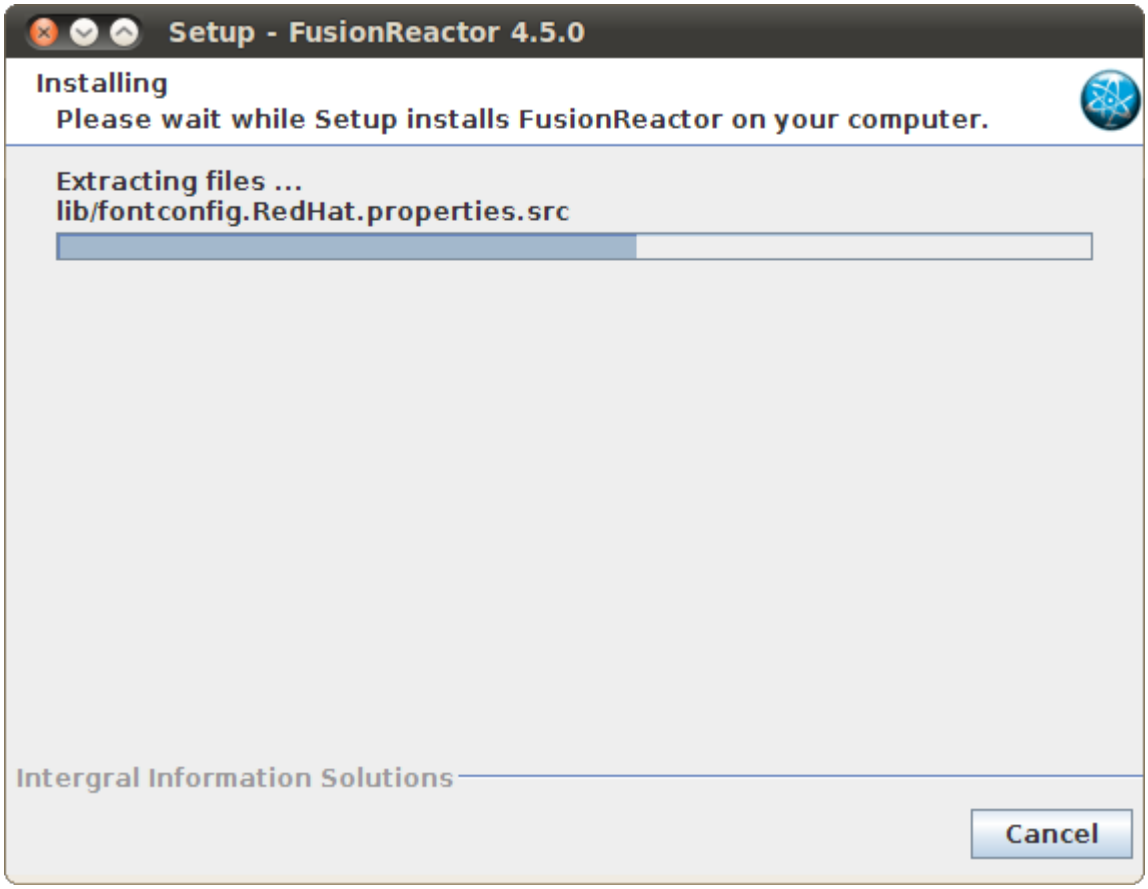


Figure 9: Installation Progress

Figure 9: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 10: Configure Instances Screen

Figure 10: After the installation has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. To configure now launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Server Scan Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close"* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 8.



Figure 11: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 11: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 8. This is the last screen in the setup, clicking Finish will close the window.

Updating FusionReactor 3.x on Linux

Executables

Executable	System Architecture
FusionReactor_linux_4_5_0.sh	32-bit
FusionReactor_linux_4_5_0_x64.sh	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

You can run the FusionReactor 4.5.x Setup to update an already installed version of FusionReactor 3.x.

To start the update, login as user *root* and run the appropriate executable as described in the above section **Exe cutables.**

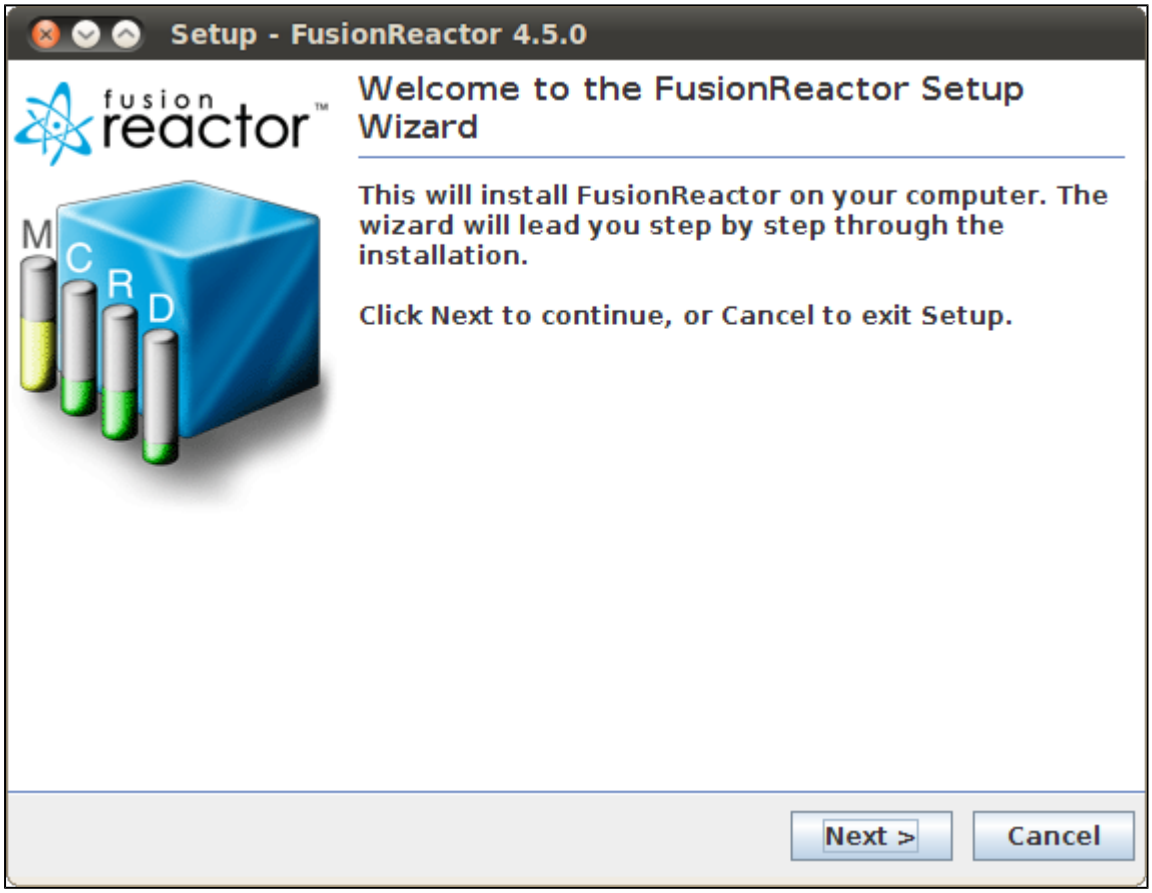


Figure 1: Welcome screen

Figure 1: The first step in the installation wizard is the welcome screen as shown in Figure 1. Clicking the Next button will proceed with the update.

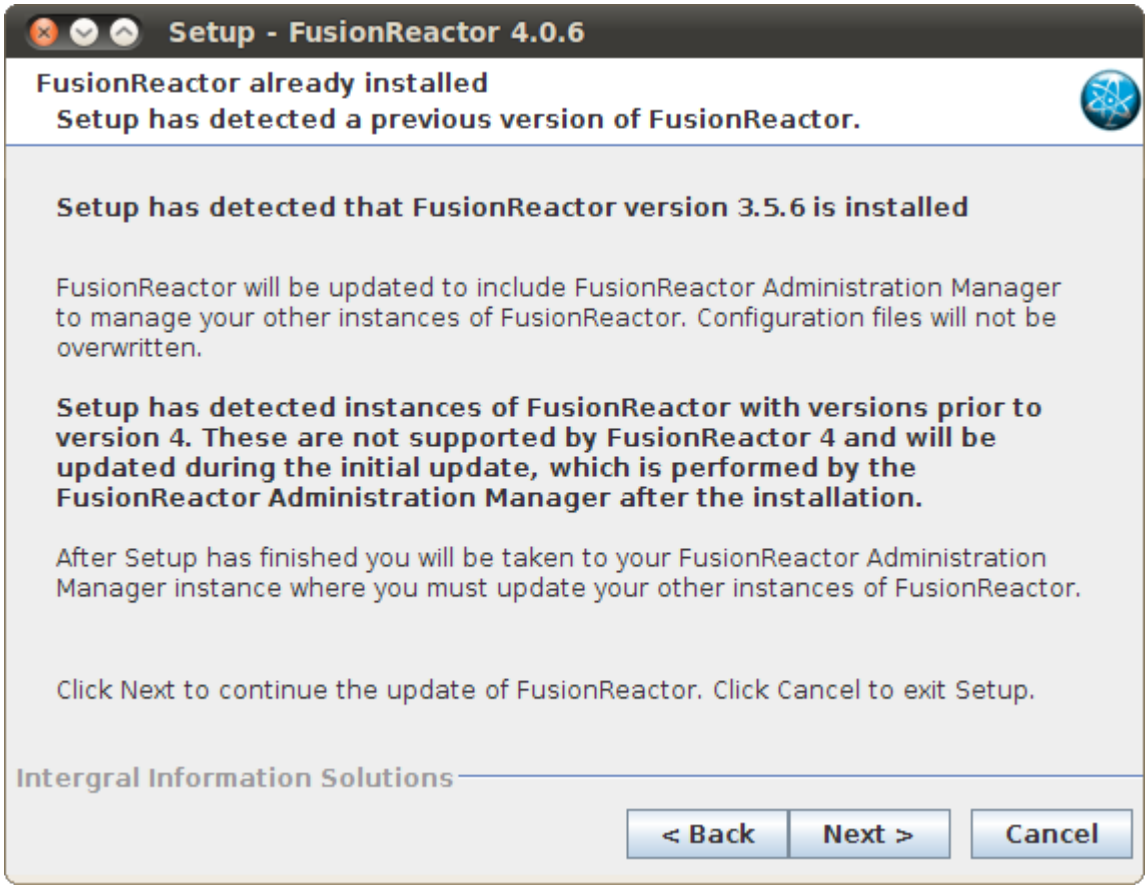


Figure 2: Update the Current Version

Figure 2: If there is already a version of FusionReactor 3.x on your computer that has been installed with the FusionReactor 3.x Setup, you will see a screen similar to the one shown. Carefully read the information and click Next to proceed with the update. Click Cancel to exit the installation wizard.

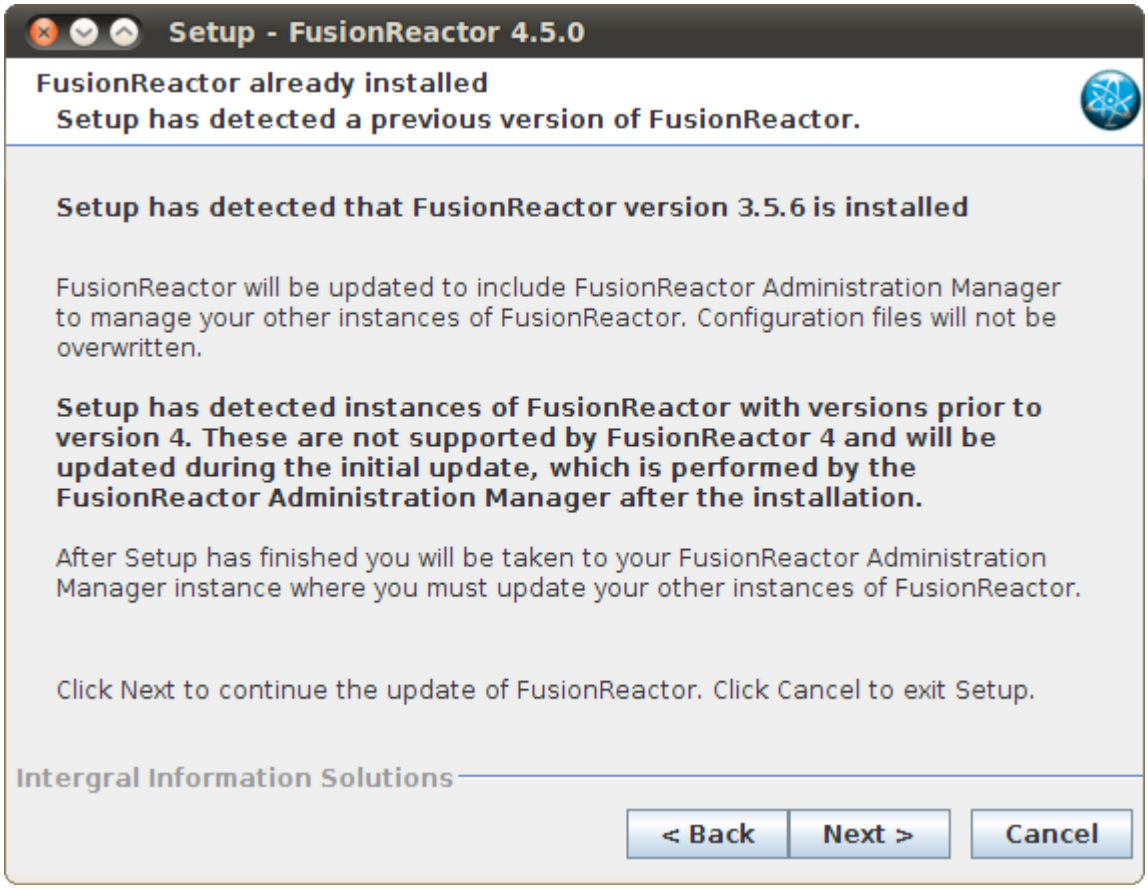


Figure 3: License Agreement

Figure 3: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.



Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the server ports (See Figure 5). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 6).

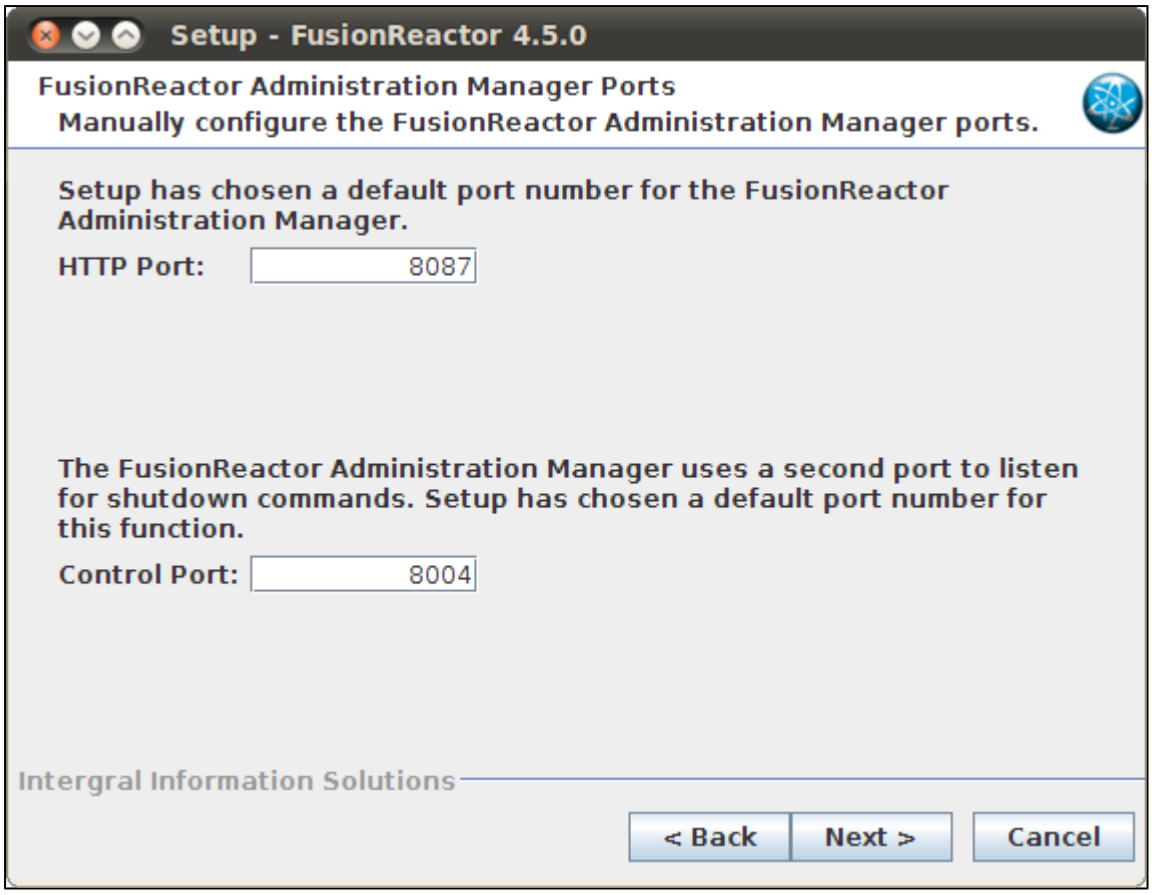


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

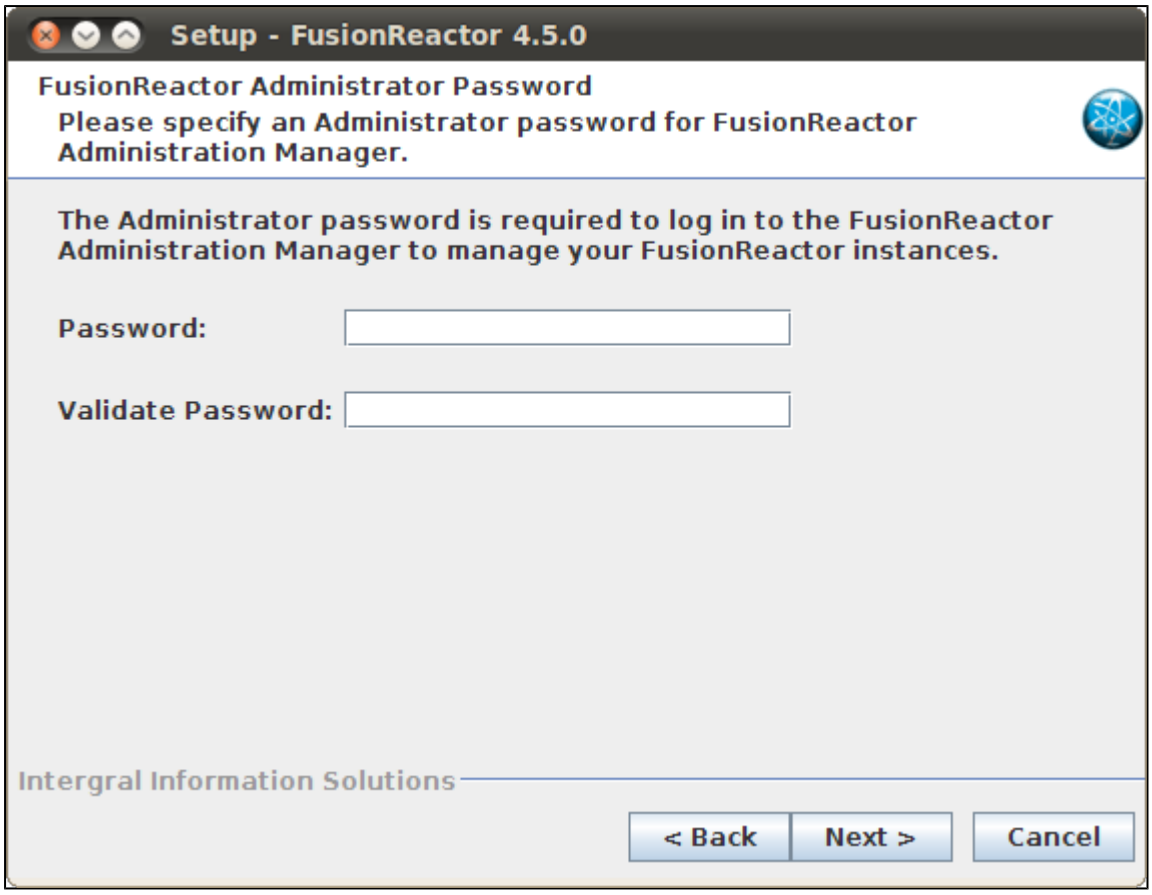


Figure 6: Administration Manager Password

Figure 6: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

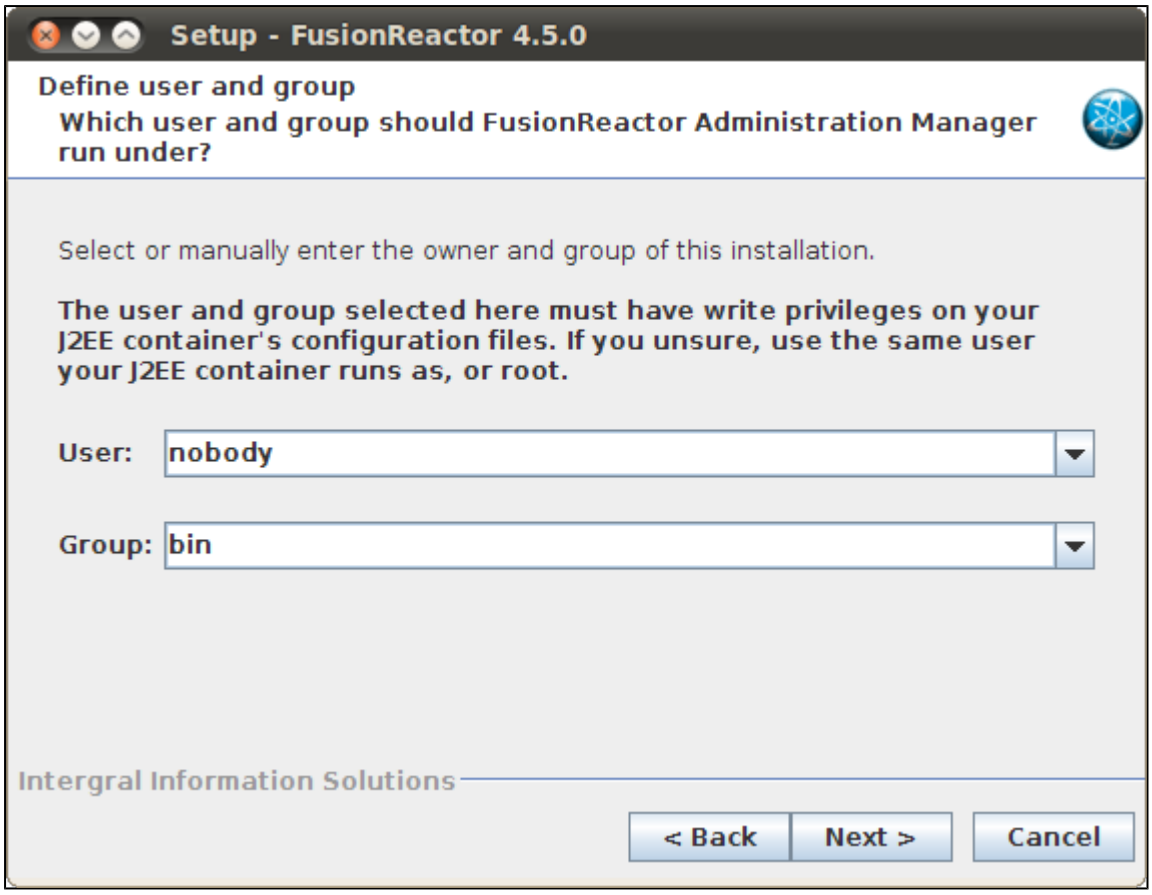


Figure 7: Choose a User and a Group

Figure 7: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM).

Please Note

FusionReactor Administration Manager (FRAM) must be installed as a user who has a real shell and can log in to the system. By default, the installer will pick **nobody/root** for this purpose, but can't presently check that this user can actually log in. If your **nobody** user does not have a shell, please change their shell to **/bin/sh** (or your preferred shell) or select another user.

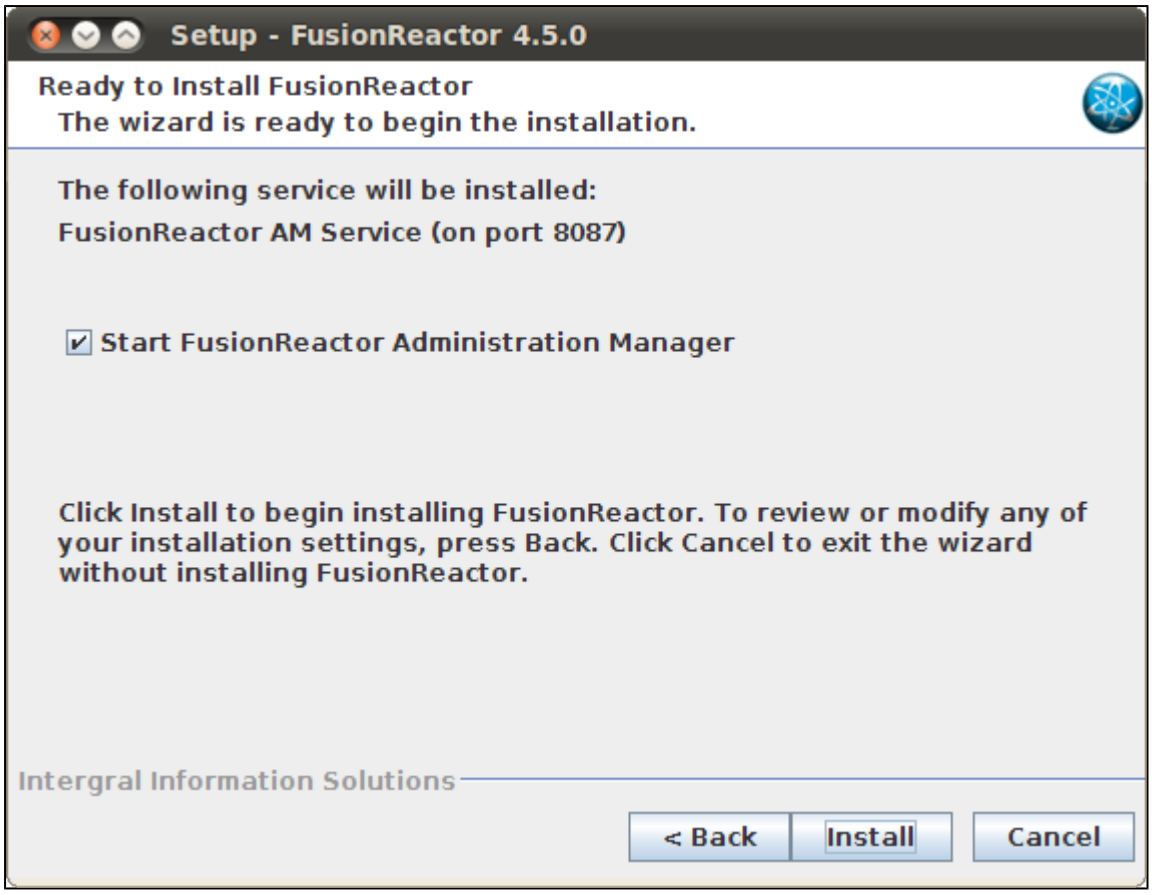


Figure 8: Ready to Install

Figure 8: Setup can automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin updating to FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

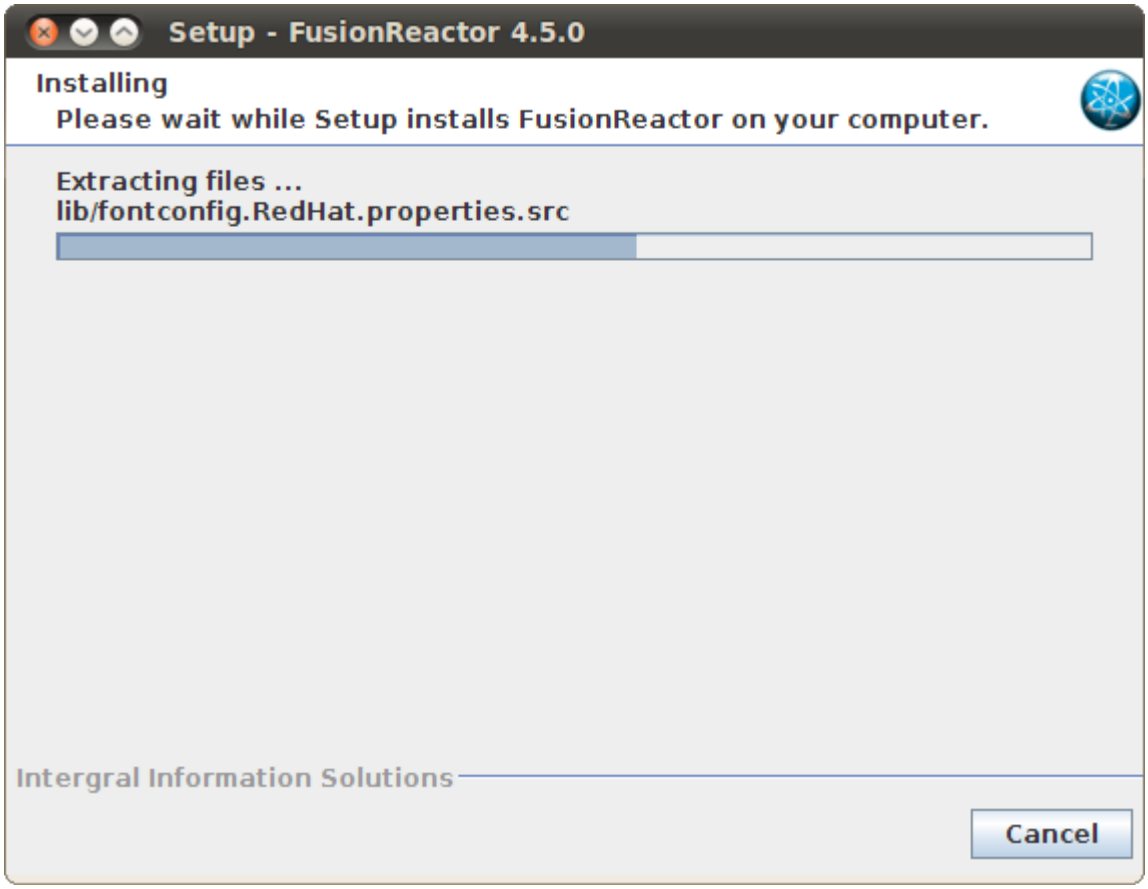


Figure 9: Installation Progress

Figure 9: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 10: Configure Instances Screen

Figure 10: After the update has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. You should now navigate to the Instance Manager located on FRAM to update your other instances of FusionReactor. Launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Update Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close'* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 8.



Figure 11: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 11: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 8. This is the last screen in the setup, clicking Finish will close the window.

Updating FusionReactor 2.0 on Linux

Upgrade Not Possible

FusionReactor 4.x.x does not support upgrades from FusionReactor 2.x.

If you currently have FusionReactor 2.x installed it is not possible to upgrade to FusionReactor 4.x.x. When running the FusionReactor 4.x.x setup you will be presented with a screen similar to that shown in Figure 1 below.

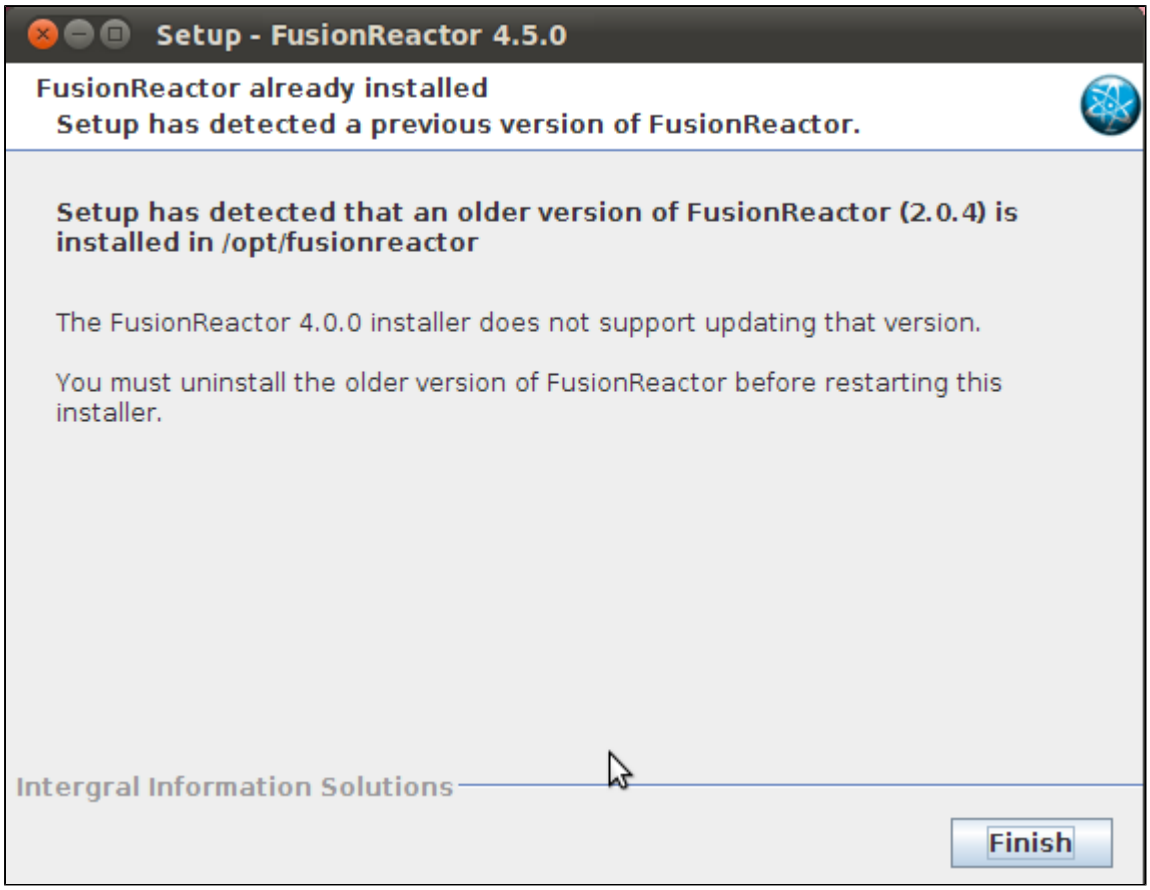


Figure 1: Update Not Supported

To install FusionReactor 4.5.x you must first uninstall FusionReactor 2.x and then run the FusionReactor 4.5.x Setup again. Please follow the guide on how to Uninstall FusionReactor 2.x on Linux

Updating FusionReactor 1.0 on Linux

> **Upgrade Not Possible**
>
> FusionReactor 4.x.x does not support upgrading from FusionReactor 1.x.

If you currently have FusionReactor 1.x installed it is not possible to upgrade to FusionReactor 4.x.x.

To install FusionReactor 4.5.x you must first uninstall FusionReactor 1.x and then run the FusionReactor 4.5.x Setup again.

Console Installation on Linux

Executables

Executable	System Architecture
FusionReactor_linux_4_5_0.sh	32-bit
FusionReactor_linux_4_5_0_x64.sh	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

The setup needs to be run as root.

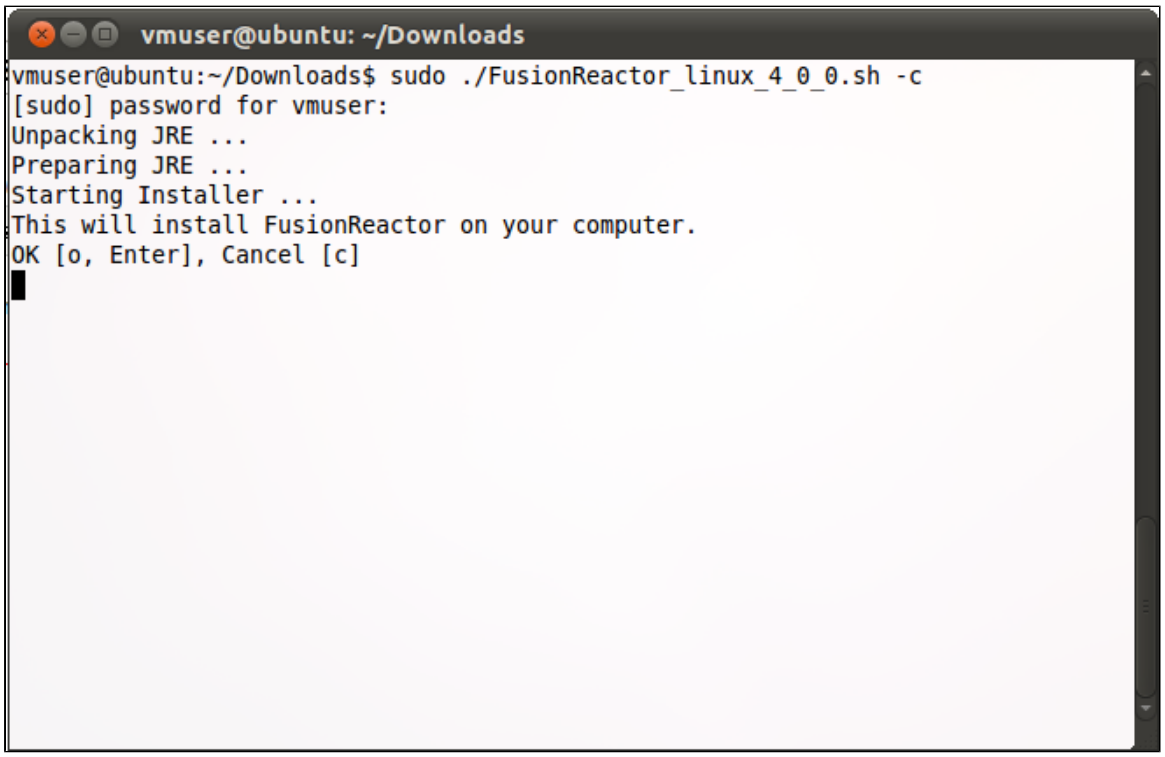


Figure 1: Start the Installation

Figure 1: To run the FusionReactor 4.5.x installation in console mode, you need to add the c switch .
For example: `sudo ./FusionReactor_linux_4_5_0.sh -c`

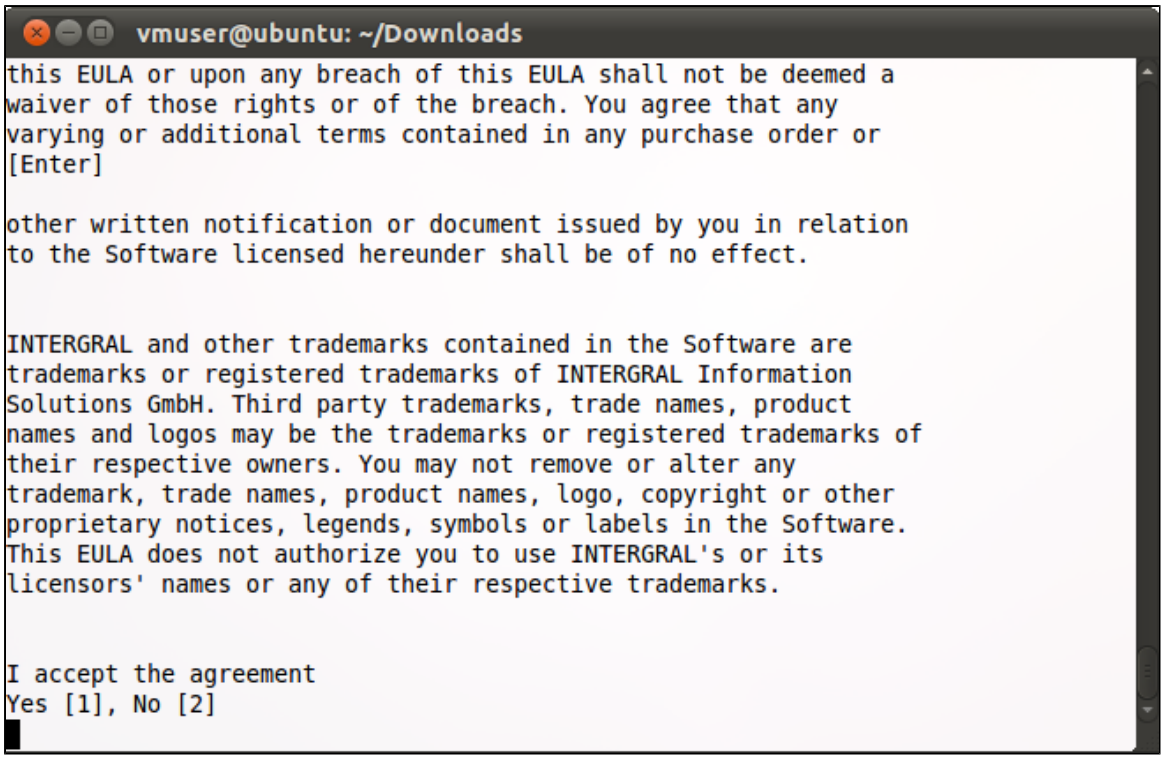


Figure 2: License Agreement

Figure 2: After starting the install you will be shown the license agreement. Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), by entering '1'.

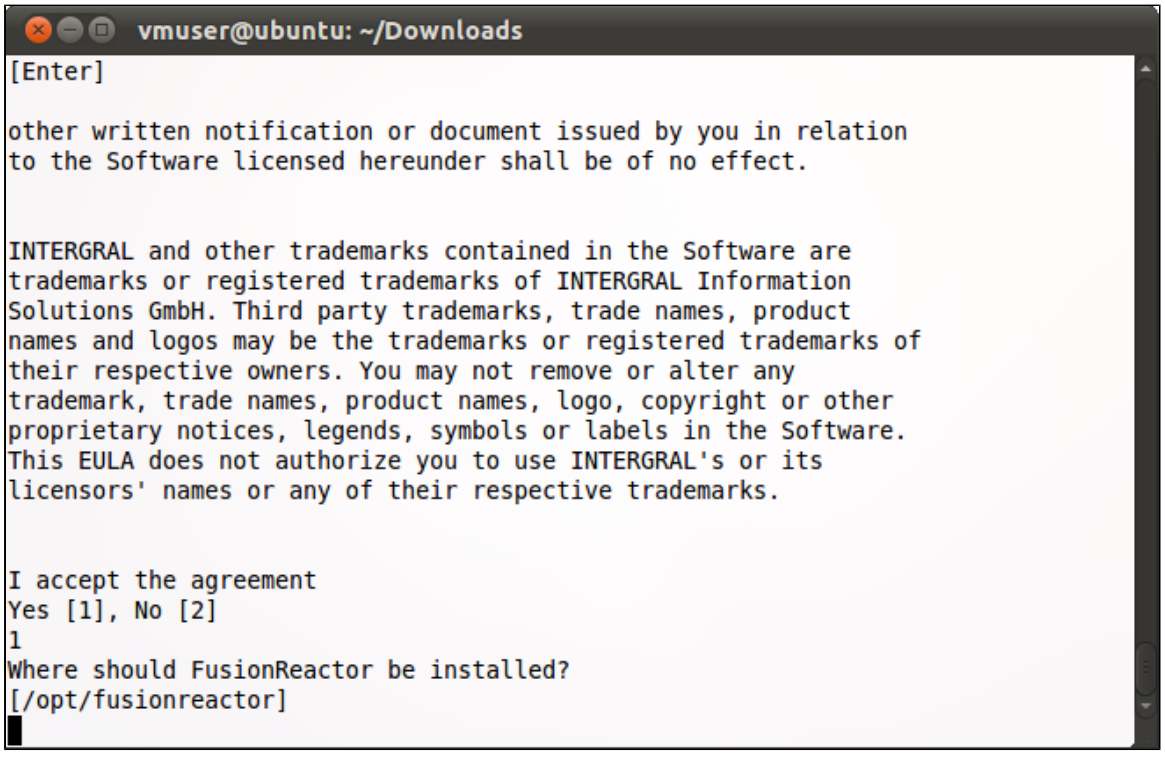


Figure 3: Installation Directory Selection

Figure 3: Now you have to select where FusionReactor is to be installed. By default (pressing enter directly), FusionReactor will be installed into `/opt/fusionreactor/`.

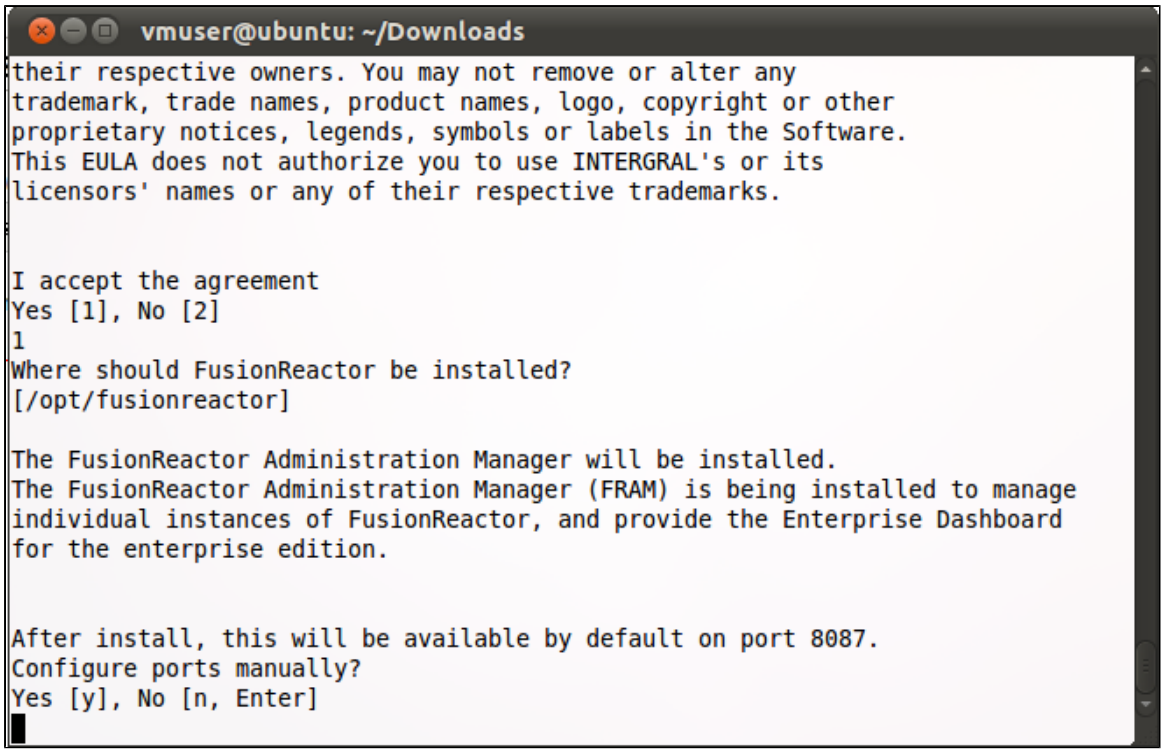


Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: You will now be asked if you want to configure the ports manually for FusionReactor Administration Manager (FRAM). Either enter 'y' to configure them (see Figure 5 onwards) or enter 'n' or just Enter on the keyboard (see Figure 7 onwards).

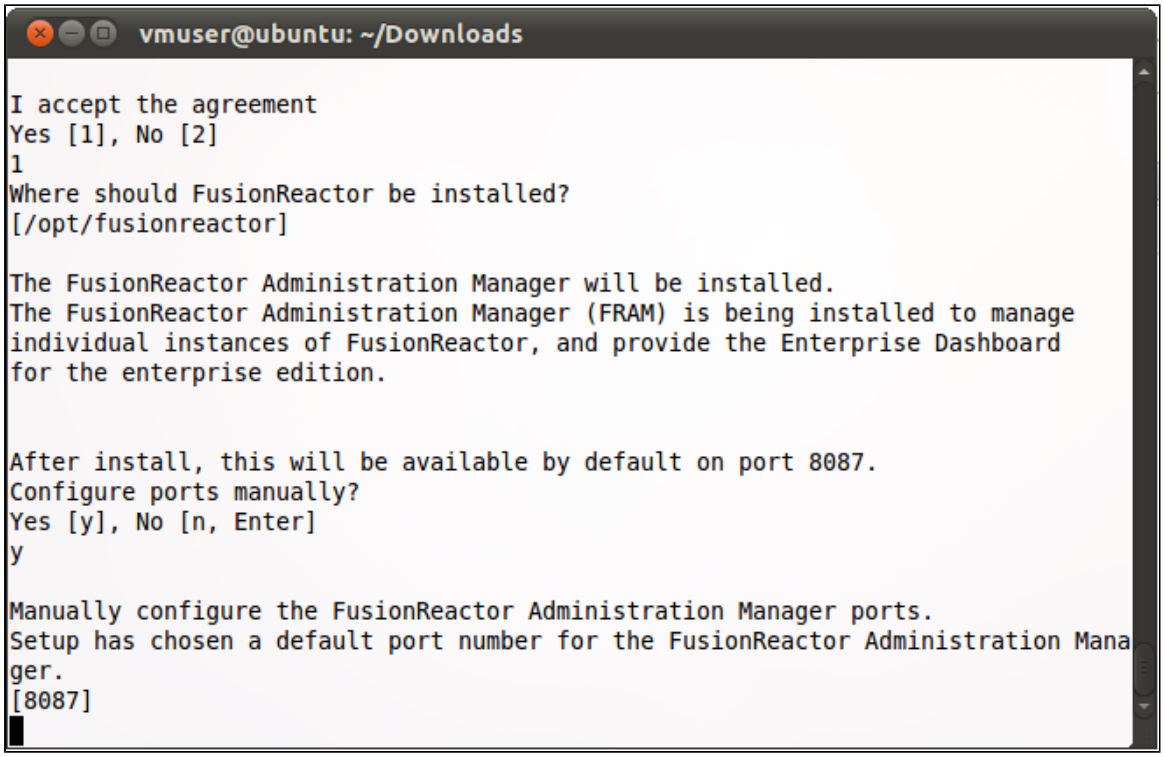


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (the port you use to view FusionReactor in a web browser). It is important not to use ports that are already in use by other applications. NOTE that the default port for FRAM is 8087.

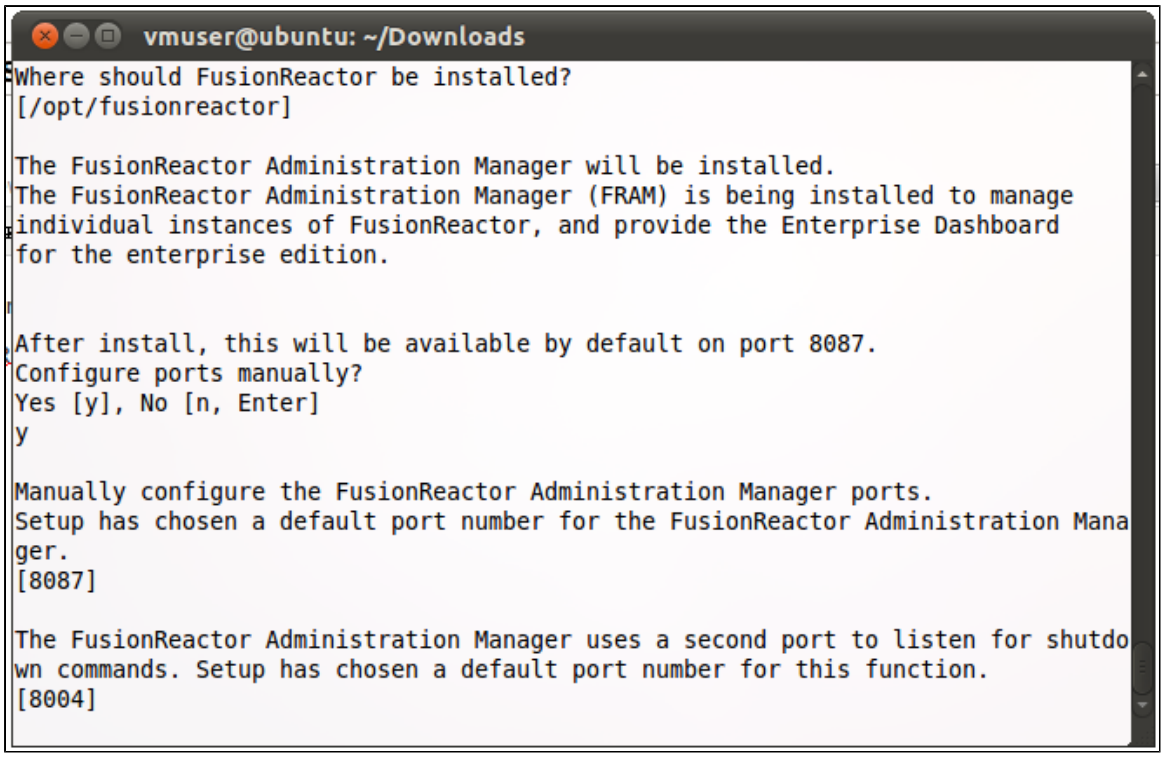


Figure 6: Manually Configuring Ports

Figure 6: You must also configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications.

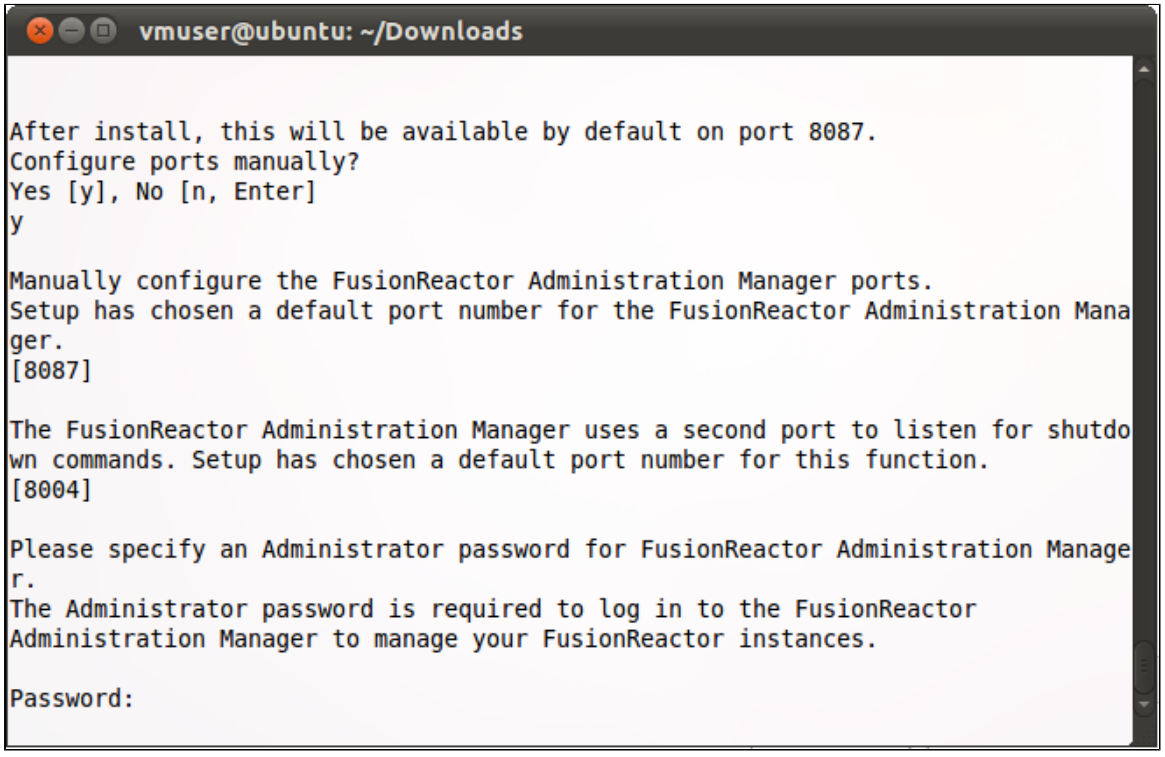


Figure 7: Administration Manager Password

Figure 7: Here you will be asked to specify a password for FRAM's Administrator account. It must be longer than 5 characters.

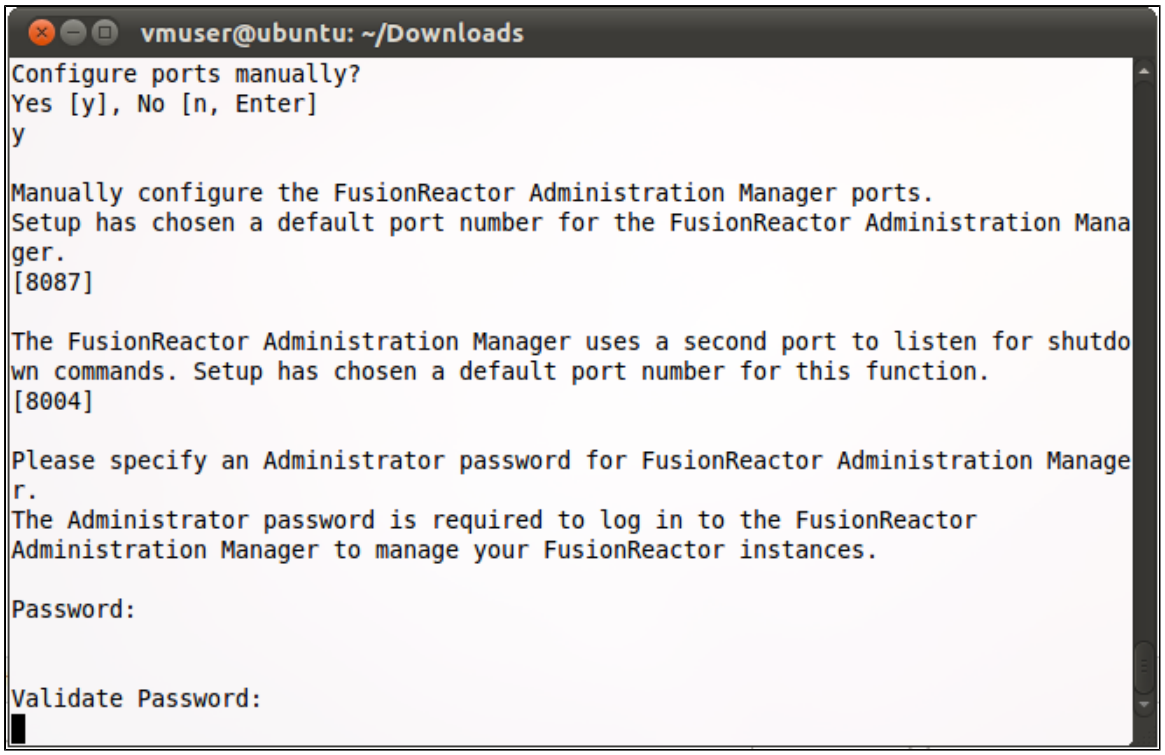


Figure 8: Administration Manager Password

Figure 8: You will be asked to validate your password.

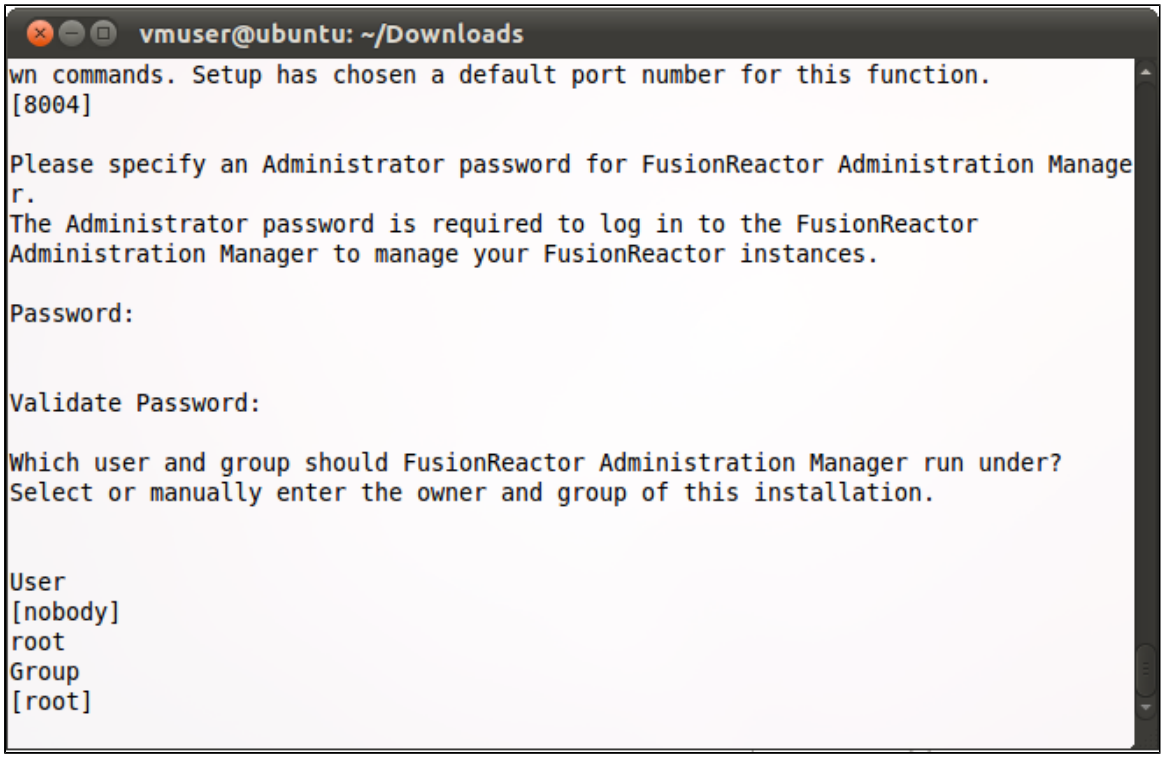


Figure 9: Choose a User and a Group

Figure 9: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM). Only that user or someone with the same or higher privileges will be allowed to run FRAM.

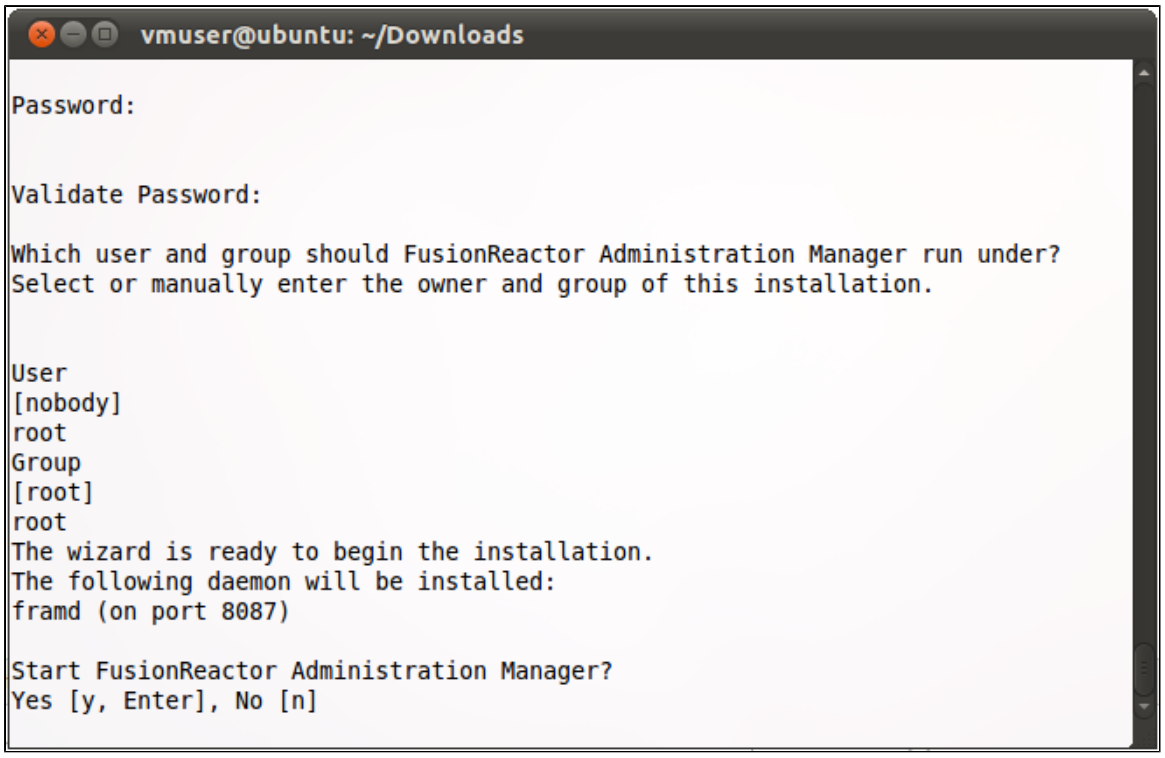


Figure 10: Ready to Install

Figure 10: Here you will be given the choice of starting up the FRAM daemon. Either entering 'y' or pressing enter will start the daemon while entering 'n' will not start the daemon. FRAM can be started manually after the installation has finished.

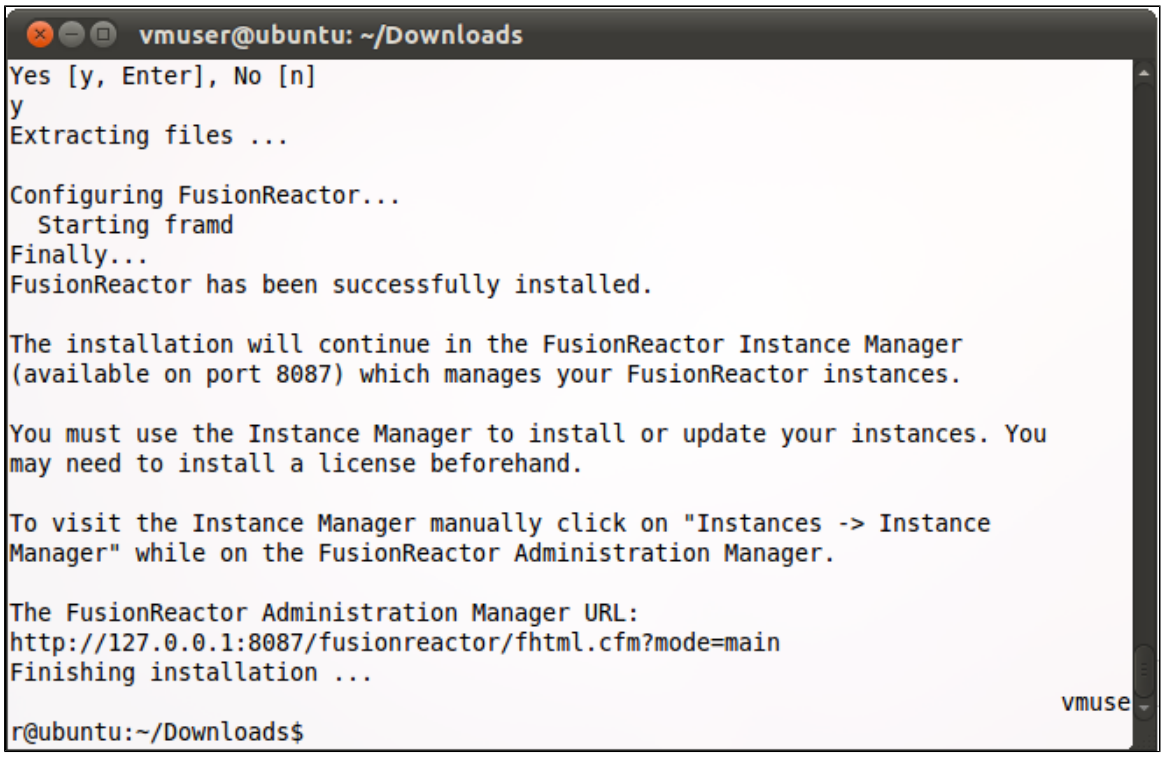


Figure 11: Finish Screen

Figure 11: You have successfully installed FusionReactor. Open the URL to go to FRAM.

Updating FusionReactor 4.0.x on Linux

Executables

Executable	System Architecture
FusionReactor_linux_4_5_0.sh	32-bit
FusionReactor_linux_4_5_0_x64.sh	64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

You can run the FusionReactor 4.5.x Setup to update an already installed version of FusionReactor 4.0.x

To start the update, login as user *Administrator* and run the appropriate executable as described in the above section **Executables.**

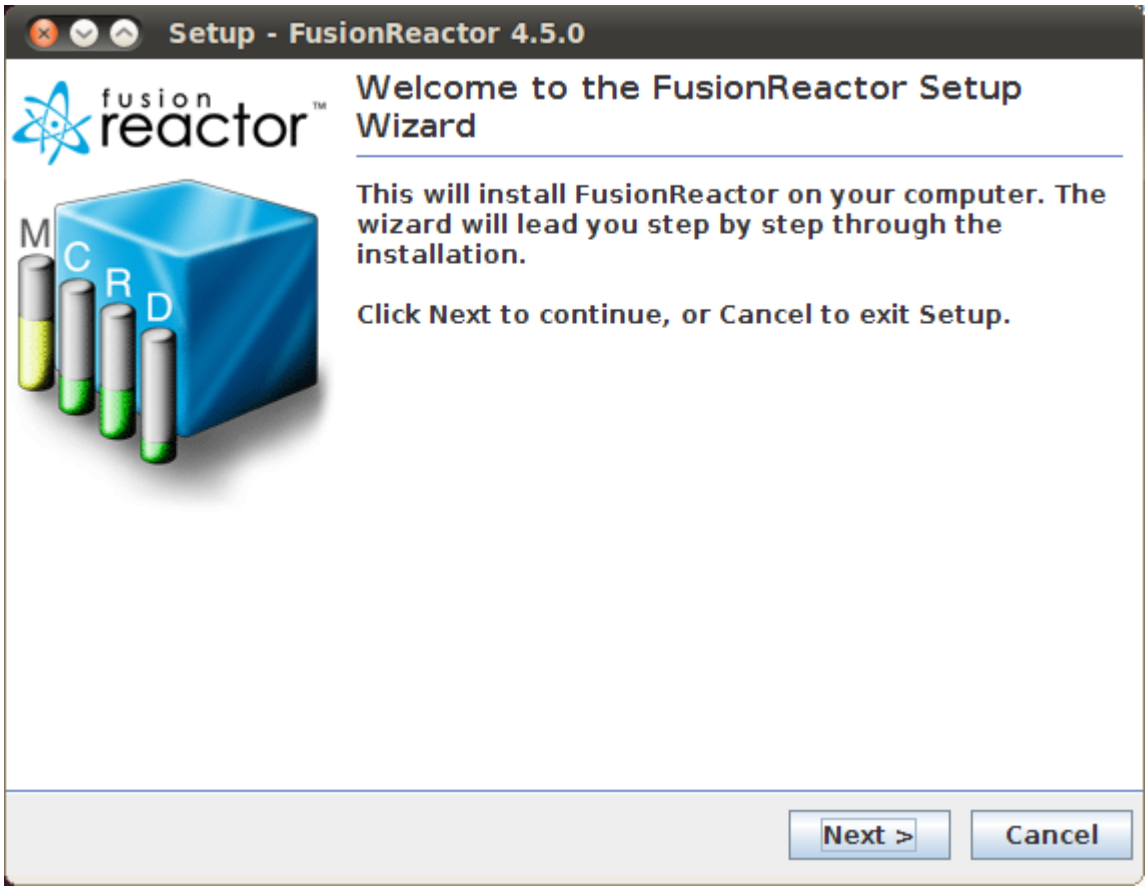


Figure 1: Welcome Screen

Figure 1: The first step in the installation wizard is the welcome screen as shown in Figure 1. Clicking the Next button will proceed with the update.

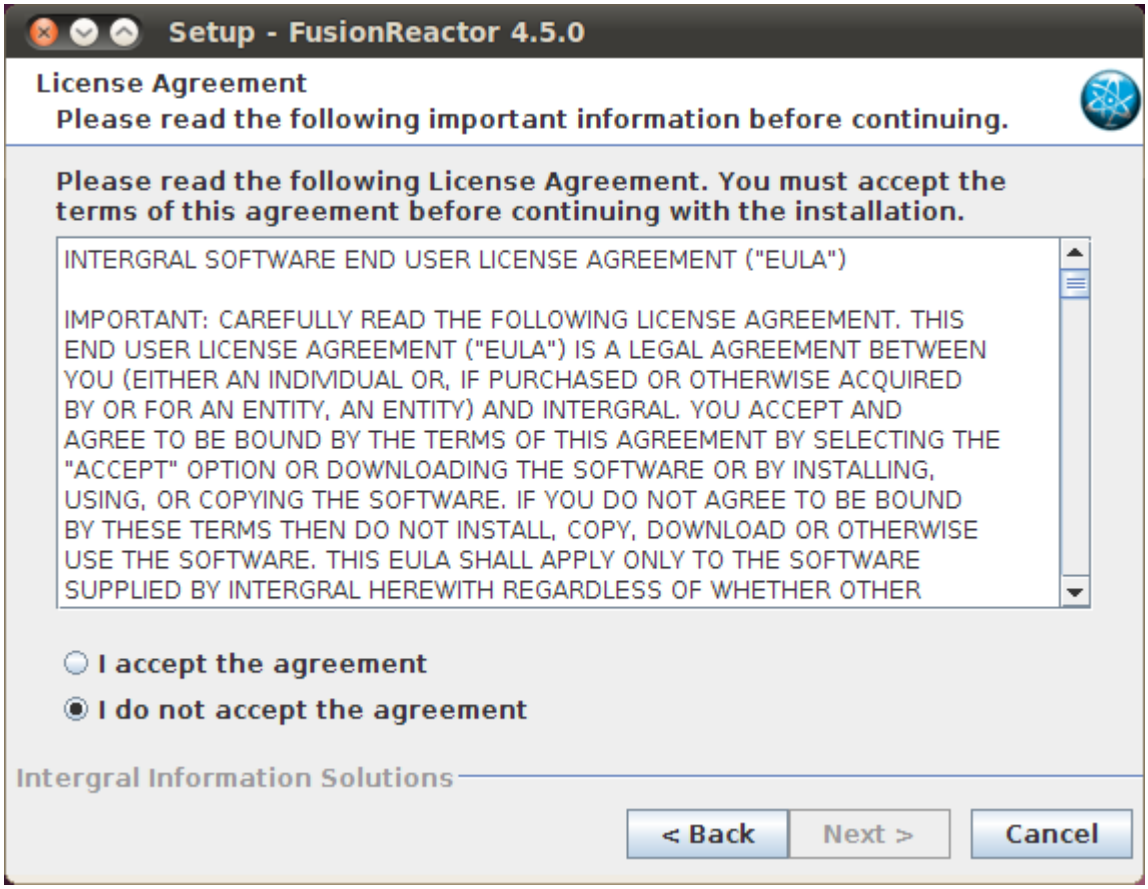


Figure 2: License Agreement

Figure 2: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.

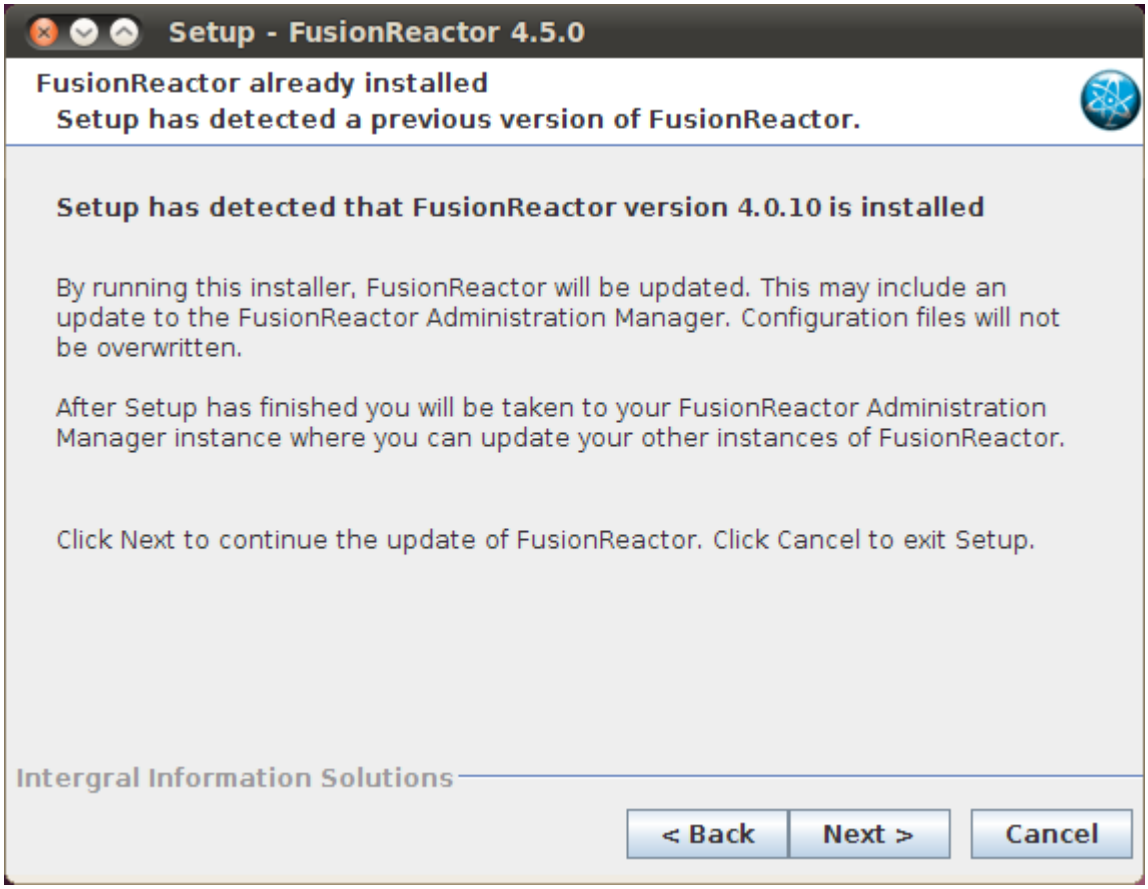


Figure 3: Update the Current Version

Figure 3: If there is already a version of FusionReactor 4.0.x on your computer that has been installed with the FusionReactor 4.0.x Setup, you will see a screen similar to the one shown. Carefully read the information and click Next to proceed with the update. Click Cancel to exit the installation wizard.

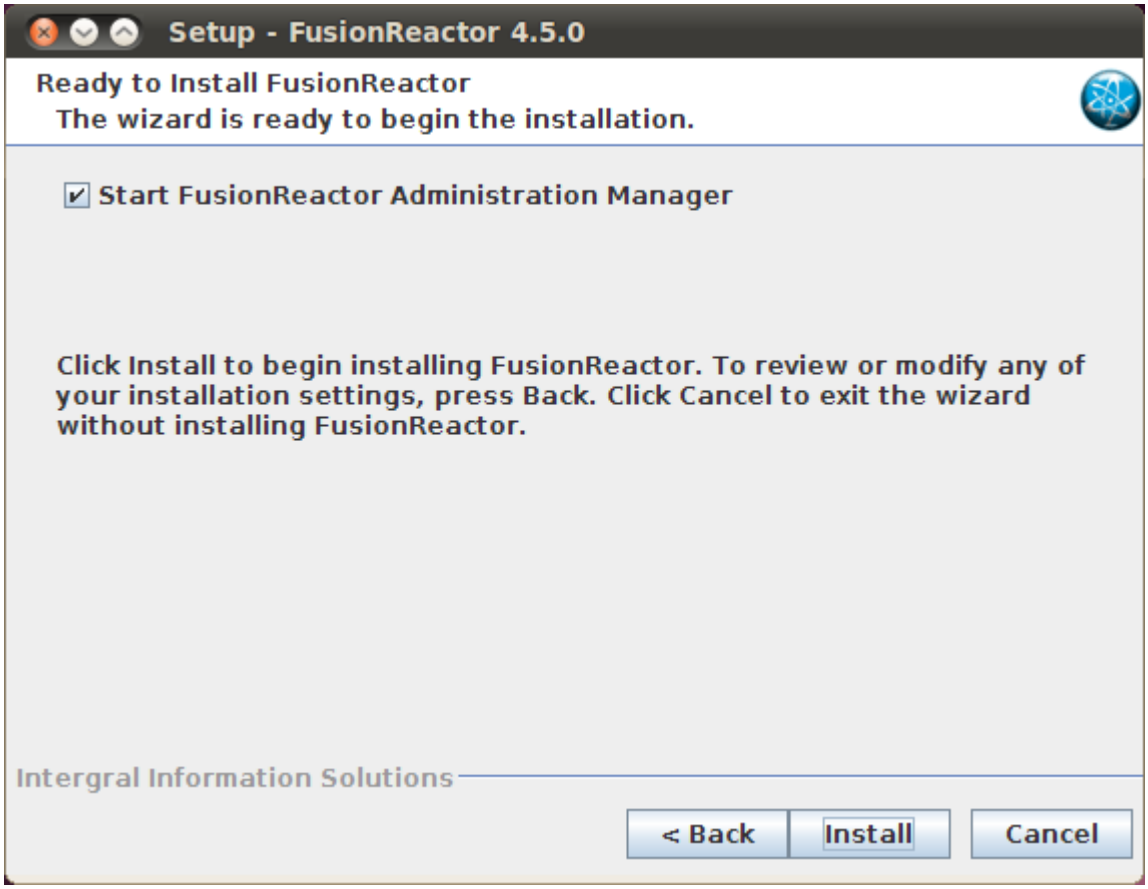


Figure 4: Ready to Install

The FusionReactor Setup can create a shortcut on the Desktop to FRAM, if you do not want Setup to perform this action uncheck the option *'Create Desktop Shortcut to FusionReactor Administration Manager".*

Setup can also automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin updating to FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

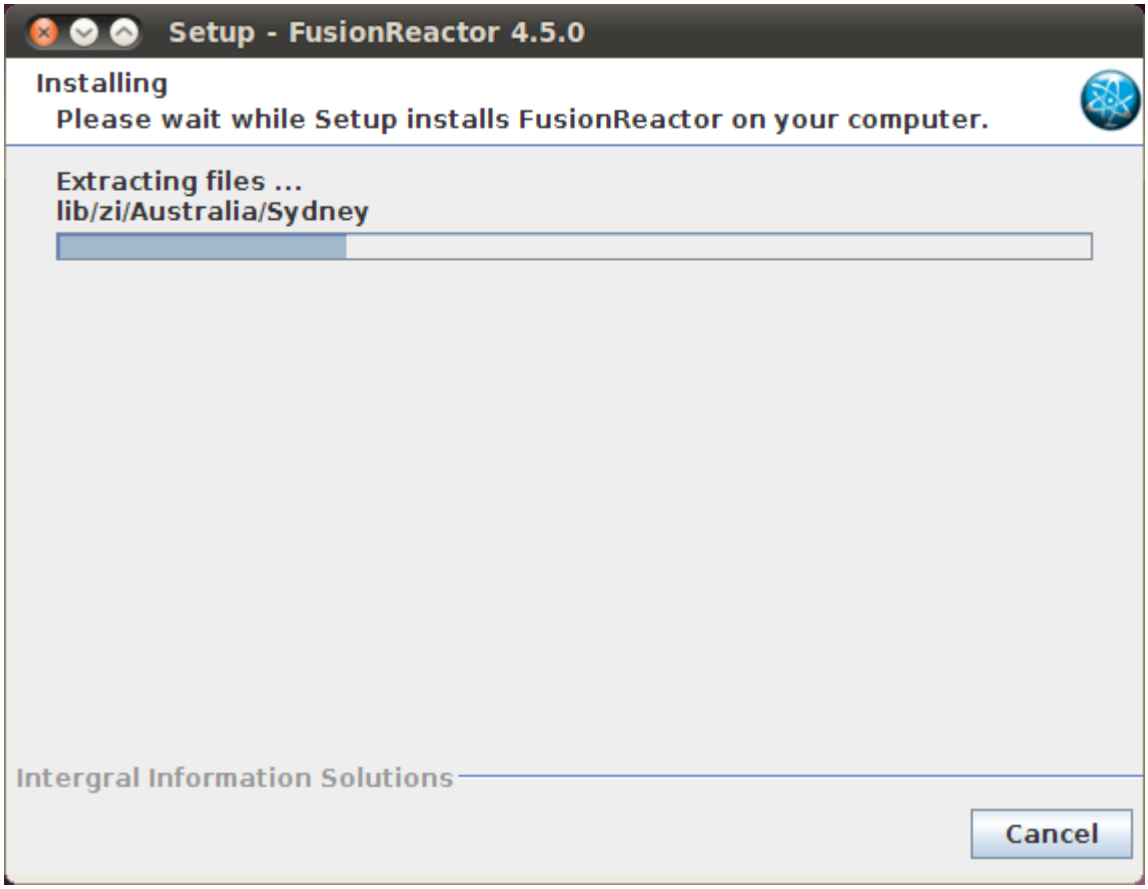


Figure 5: Installation Progress

Figure 5: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 6: Configure Instances Screen

Figure 6: After the update has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. You should now navigate to the Instance Manager located on FRAM to update your other instances of FusionReactor. Launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Update Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close'* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 4.

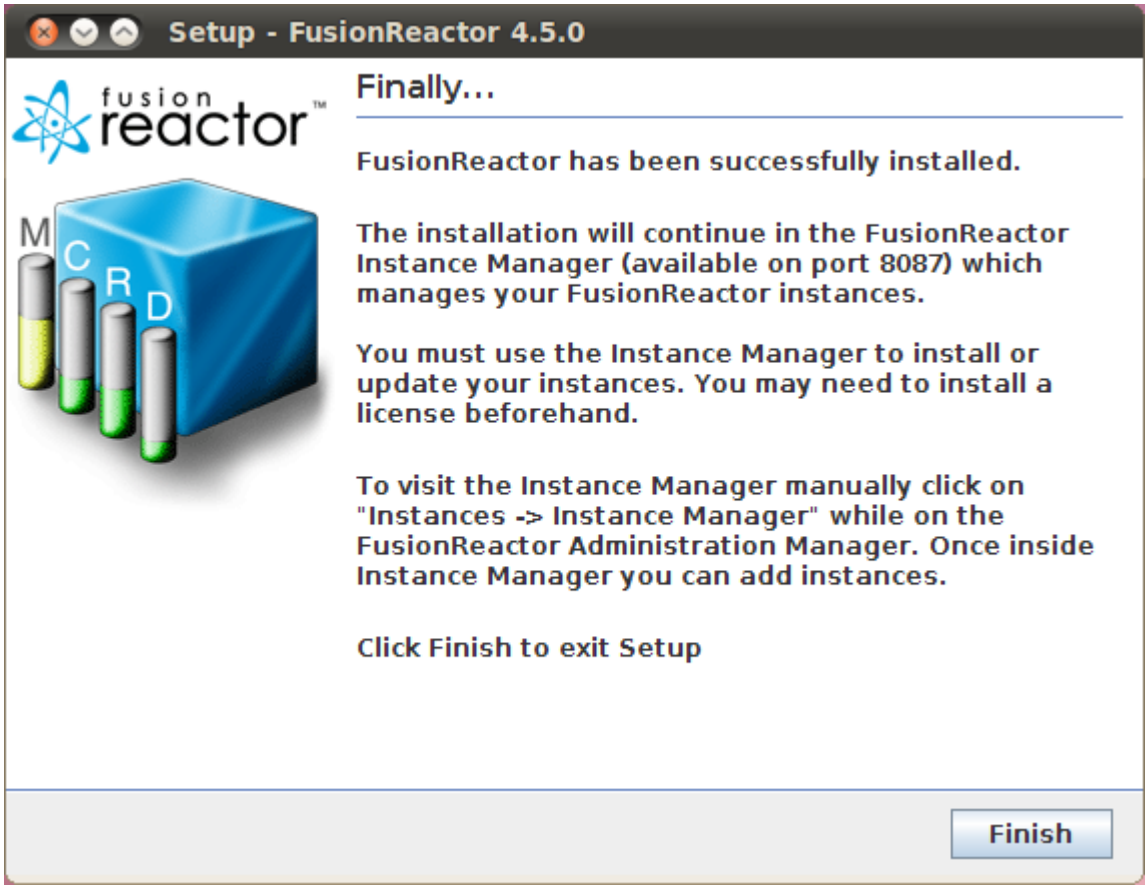


Figure 7 Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 7: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 4. This is the last screen in the setup, clicking Finish will close the window.

Uninstalling FusionReactor on Linux

To completely remove FusionReactor from your system, you can either double click on the FusionReactor Uninstaller icon in the FusionReactor directory or you can run the uninstall command located in the FusionReactor directory in a shell window. Both methods require you to be logged in as the root user.

The Uninstaller will completely remove the FusionReactor directory including all of the generated data – **please take care that you save any files you want to keep before running the uninstaller.**

> The uninstaller will also remove any installed FusionReactor instances, during this process you will be asked to restart any running servers. Please note that not stopping the servers means that FusionReactor and all of its components will not be completely removed. This may cause issues on these servers.



Figure 1: FusionReactor Uninstall

Figure 1: Once the Uninstaller has been started, read the text that asks you if you want to remove FusionReactor and all it's components. Click Next to continue on with the uninstall or click Cancel to exit the uninstaller wizard.

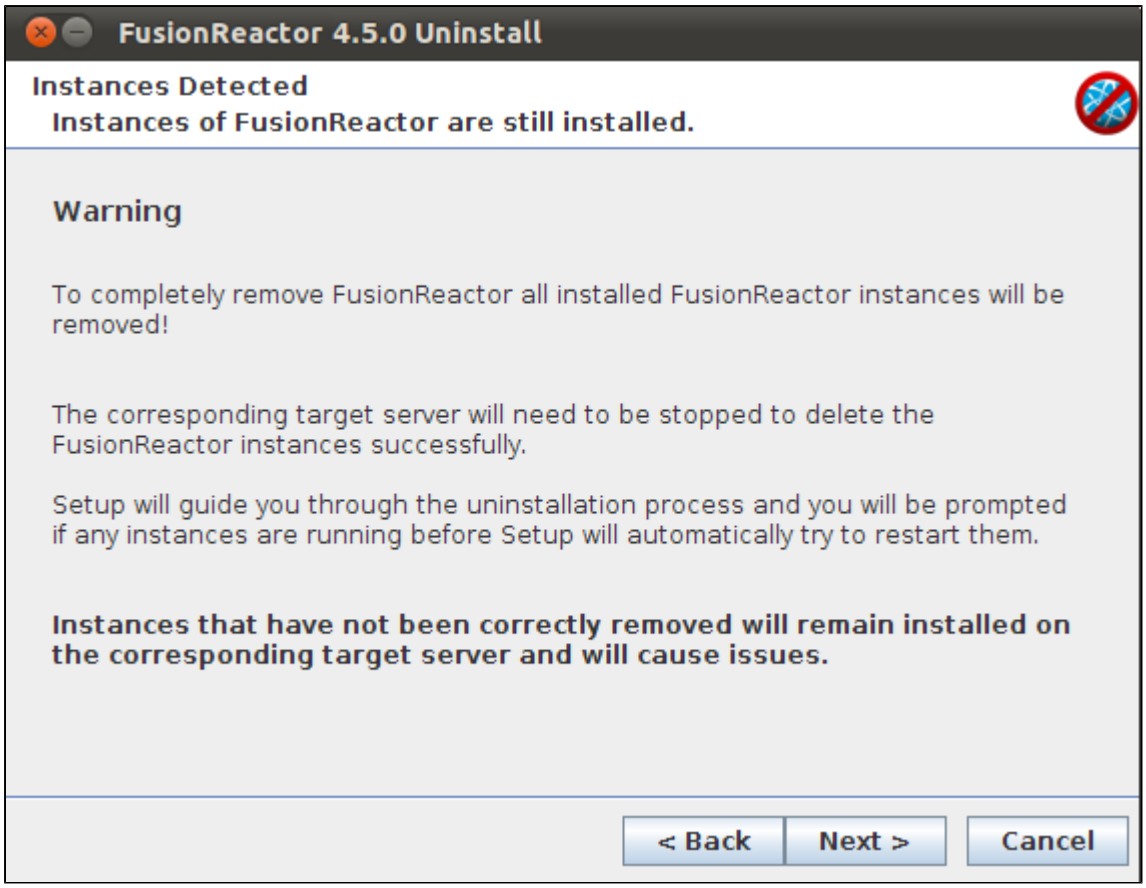


Figure 2: Uninstall FusionReactor Instances Warning

Figure 2: If you have other FusionReactor instances installed a warning message will be displayed to highlight the effects of removing FusionReactor. Carefully read through this. Once you have read it and you have understood it, please **click Next to uninstall all of the installed FusionReactor instances as well as FusionReactor.** Otherwise, click Cancel to stop the uninstall.

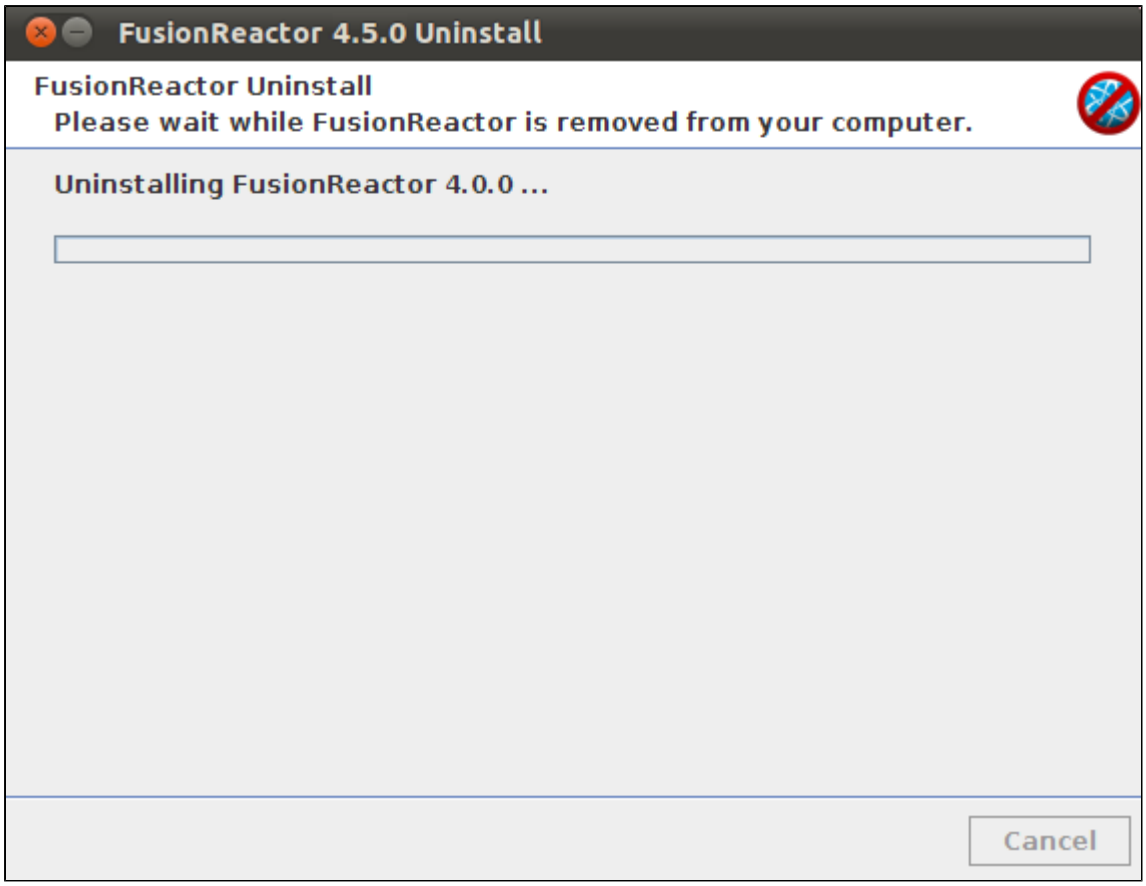


Figure 3: Uninstall Progress Bar

Figure 3: Once you have clicked Next, the FusionReactor uninstaller will begin to uninstall FusionReactor and all of the installed FusionReactor instances.

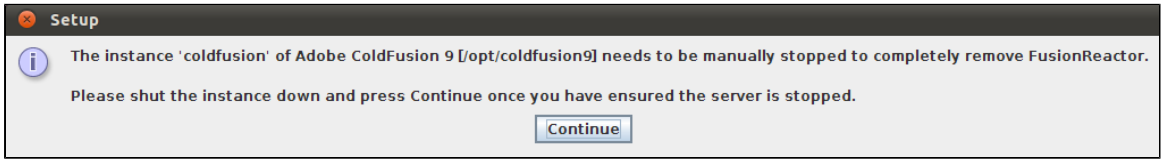


Figure 4: Stopping a Server Instance

Figure 4: If there is a FusionReactor instance that is running, the server container for this needs to be stopped in order to successfully remove the FusionReactor Instance.

Stop the server on which the FusionReactor instance is installed and press Continue once ensured the server is shut down, or press Continue if you do not wish to stop the server.

> As stated at the top of this page, not stopping the server container means that FusionReactor and all of its components will not be completely removed from the target server. This may cause issues on the corresponding server.

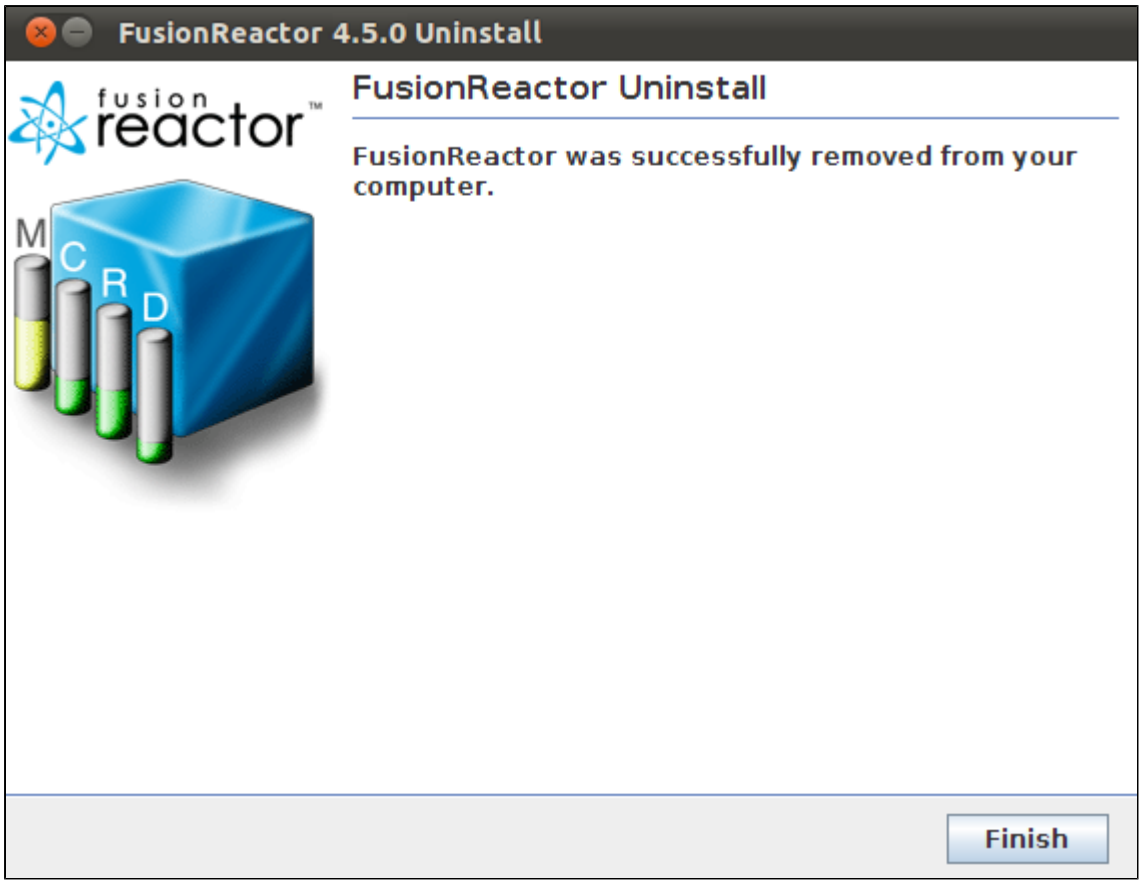


Figure 5: FusionReactor Uninstall Complete

Figure 5: When FusionReactor has finished uninstalling you will be presented with the Finish screen shown in Figure 6. Click on Finish to exit.

Overview

Welcome to the FusionReactor Installation Guide

Thank you for choosing FusionReactor to monitor your ColdFusion or J2EE server. You have reached the Intergral Documentation Center which contains all the information you will need to get up and running with FusionReactor and other Intergral products.
If you are looking for another page you can use our search feature located to the top right of every page.

These pages will describe how to use the FusionReactor Setup to install FusionReactor on Windows, Linux, Mac and Solaris. If you have not already done so you can download FusionReactor by clicking the download link to the right, if you are unsure if you can run the latest version on your server or would like to know if your operating system is supported please visit our requirements page.

Please continue to select the appropriate platform and we will guide you through the installation process.

Download

Requirements

Please Select Your Platform









The FusionReactor Setup will install the FusionReactor Administration Manager (FRAM). FRAM provides a completely isolated instance of FusionReactor which is just for monitoring and instance management, it will continue running even following a JVM crash. FRAM manages and keeps track of all servers (and instances) so you can quickly and easily install FusionReactor across your environment.

The Setup also allows you to update (or reinstall) an existing installation of FusionReactor 4.5.x as well as update an old version of FusionReactor that has been installed with the 3.x installer. The FusionReactor 4.5.x Setup does not support updating previous installations of FusionReactor 2 or FusionReactor 1.

Installing FusionReactor on Solaris

Installing FusionReactor 4.5.x on Solaris

Updating FusionReactor 3.x on Solaris

Updating FusionReactor 2.0 on Solaris

Updating FusionReactor 1.0 on Solaris

Console Installation on Solaris

Installing FusionReactor 4.5.x on Solaris

Executables

Executable	System Architecture
FusionReactor_solaris_sparc_4_5_0_x64.sh	SPARC 64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

To start the installation, login as user *root* and run the appropriate executable as described in the above section **Executables.**
If you can not use an X-Windows session, you must install FusionReactor through the console installation. Please follow the guide that covers Console Installation on Solaris



Figure 1: Welcome screen

Figure 1: The first step in the installation wizard is the welcome screen.

Clicking the Next button will proceed with the installation and open the License Agreement as shown in Figure 2. Clicking Cancel will exit the Setup.

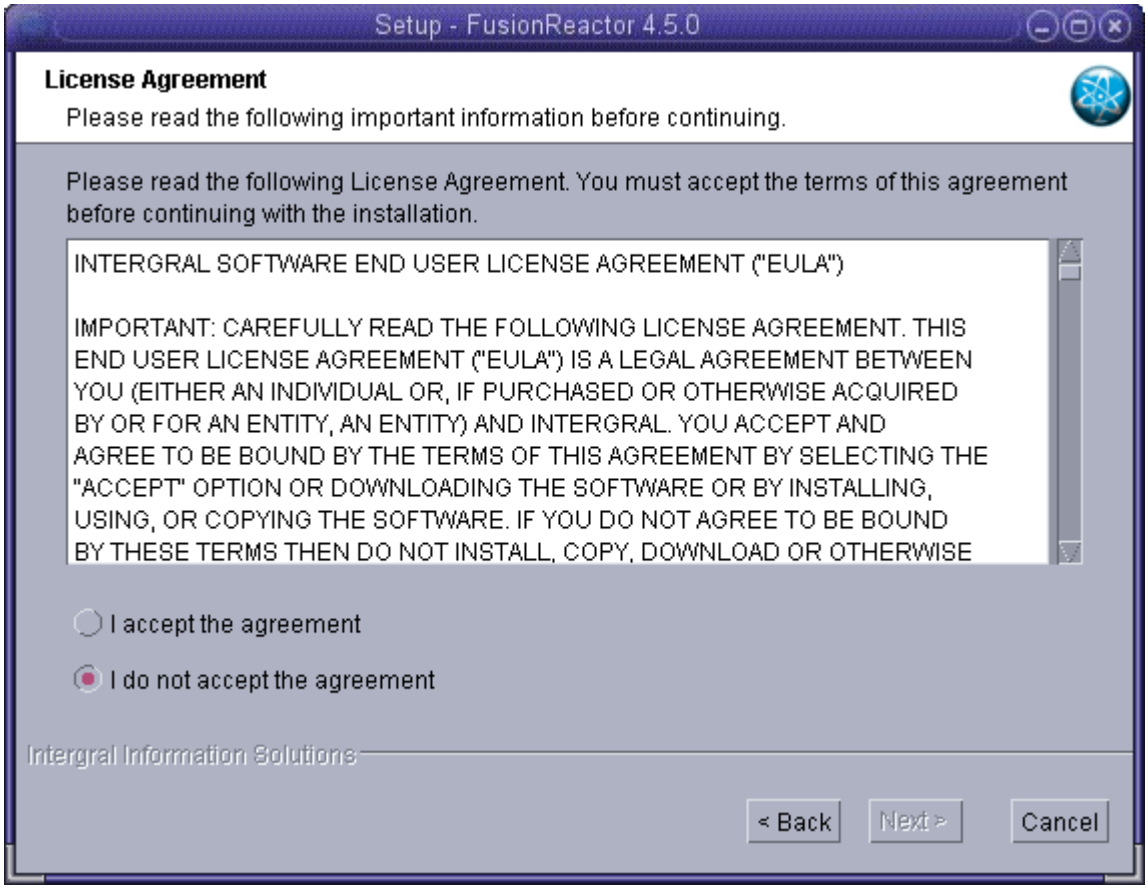


Figure 2: License Agreement

Figure 2: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.

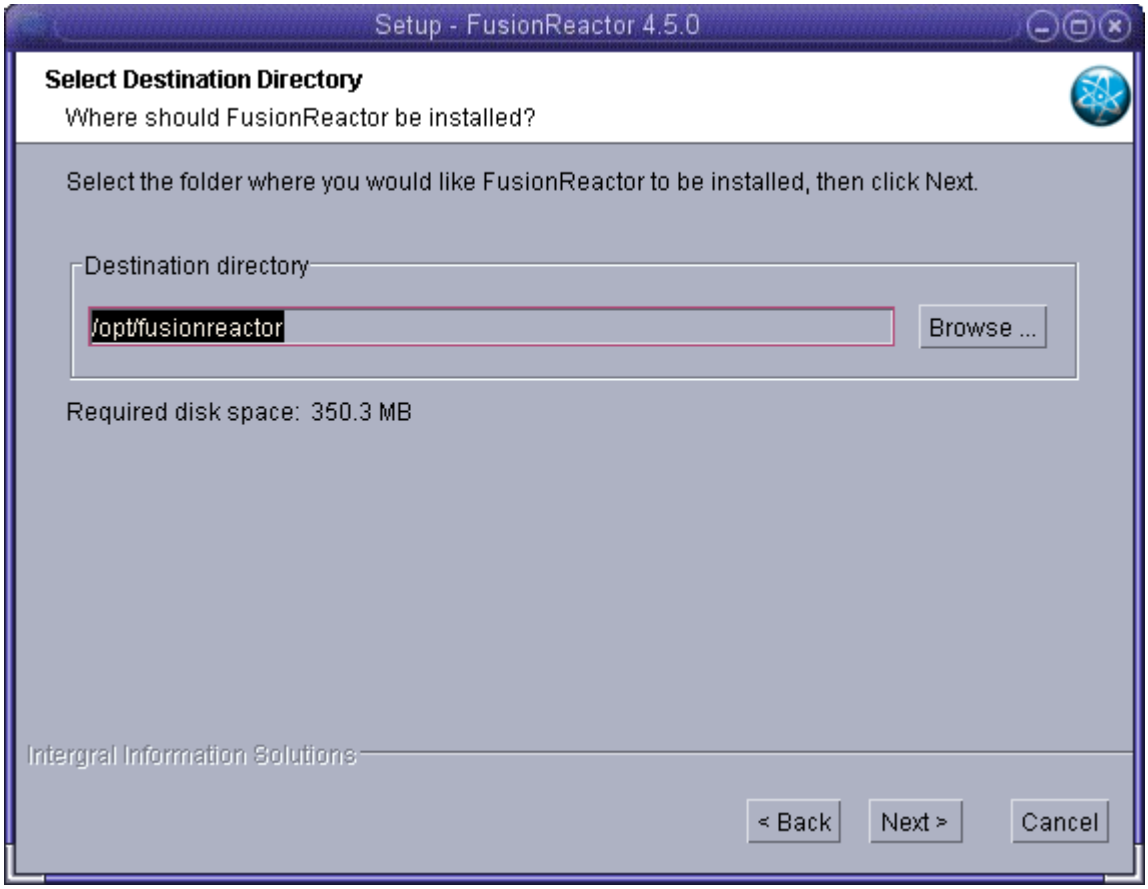


Figure 3: Installation Directory Selection

Figure 3: Now you have to select the directory in which FusionReactor should be installed. By default FusionReactor will be installed into `/opt/FusionReactor/`



Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the sever ports (See Figure 5). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 6).

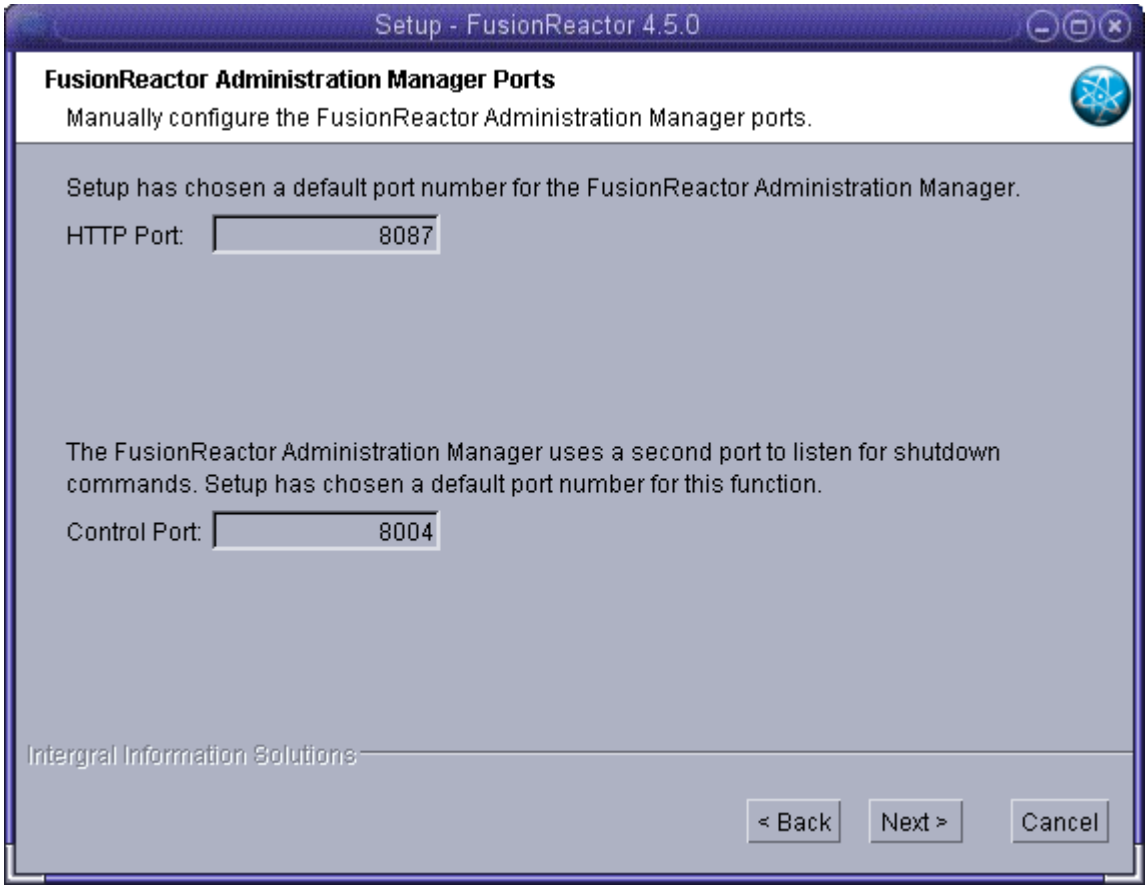


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

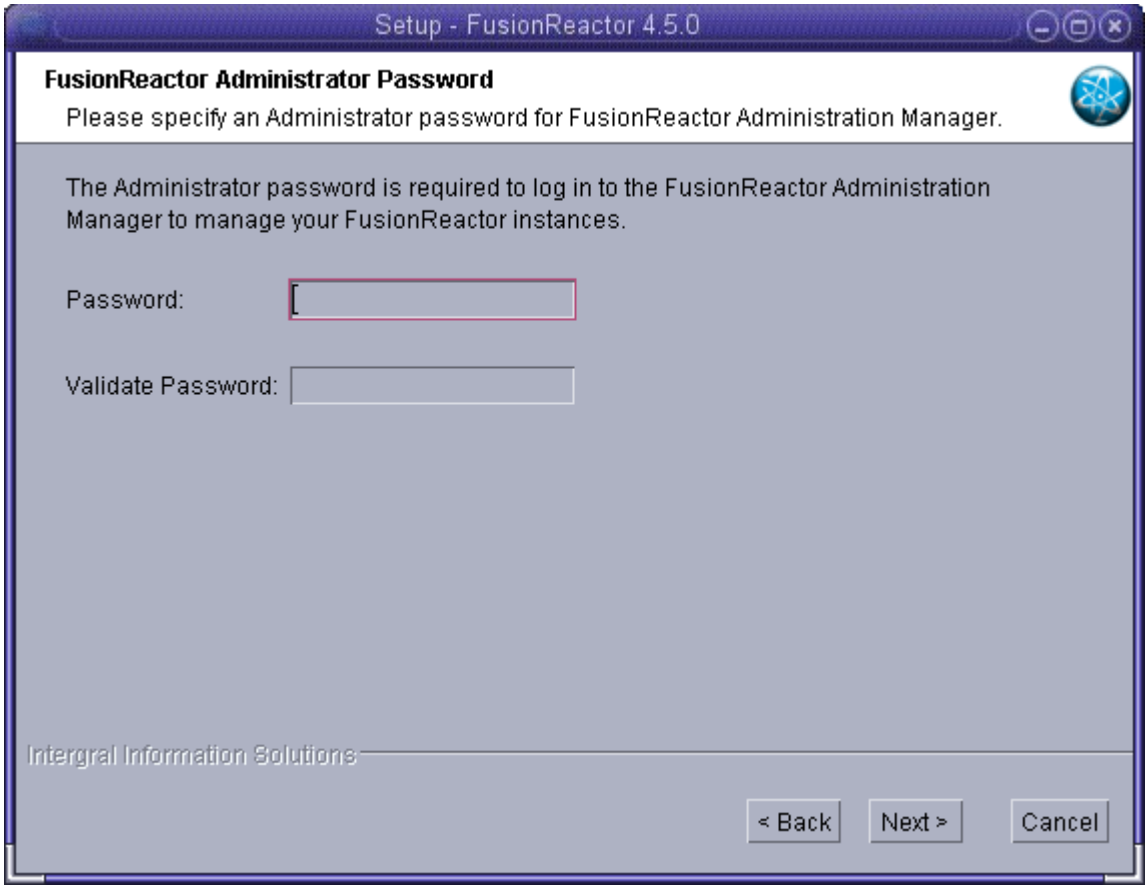


Figure 6: Administration Manager Password

Figure 6: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

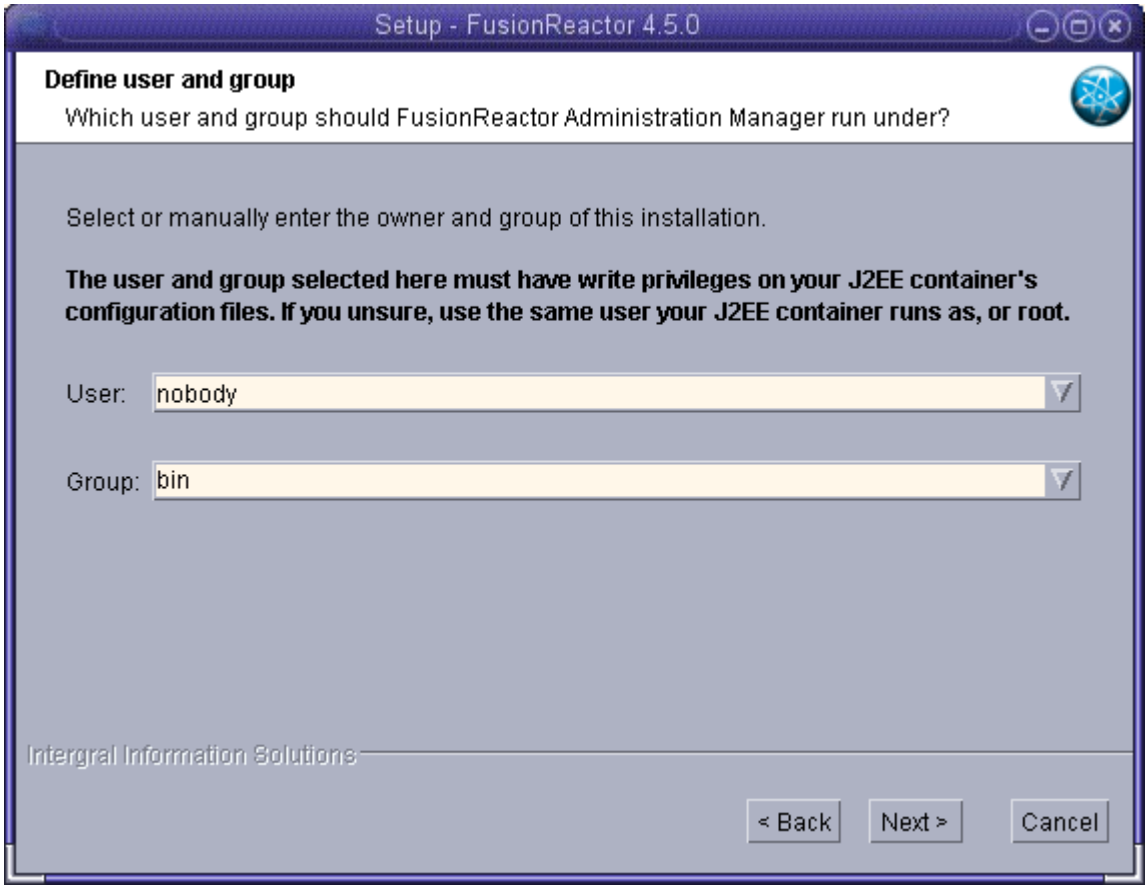


Figure 7: Choose a User and a Group

Figure 7: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM). Only that user or someone with the same or higher privileges will be allowed to run FRAM.

Please Note

FusionReactor Administration Manager (FRAM) must be installed as a user who has a real shell and can log in to the system. By default, the installer will pick **nobody/root** for this purpose, but can't presently check that this user can actually log in. If your **nobody** user does not have a shell, please change their shell to **/bin/sh** (or your preferred shell) or select another user.

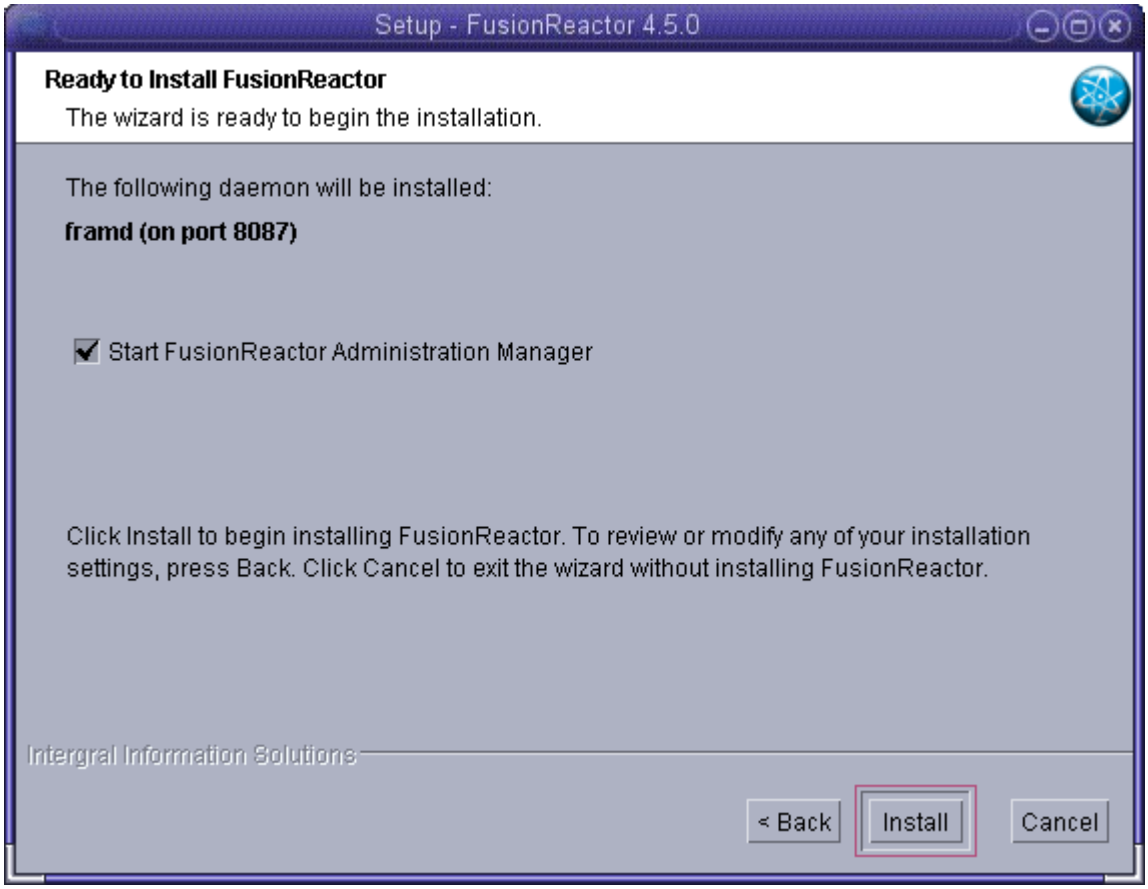


Figure 8: Ready to Install

Figure 8: Setup can automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin installing FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

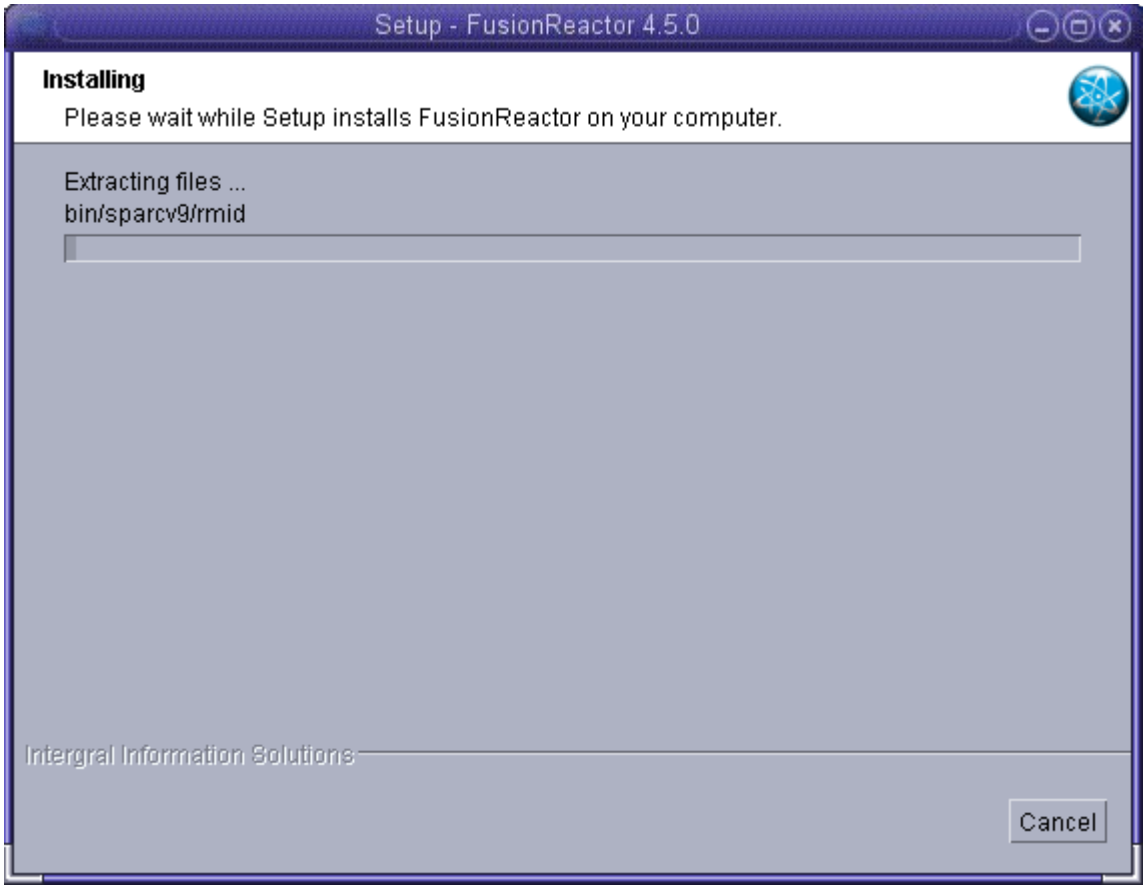


Figure 9: Installation Progress

Figure 9: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 10: Configure Instances Screen

Figure 10: After the installation has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. To configure now launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Server Scan Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close"* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 8.

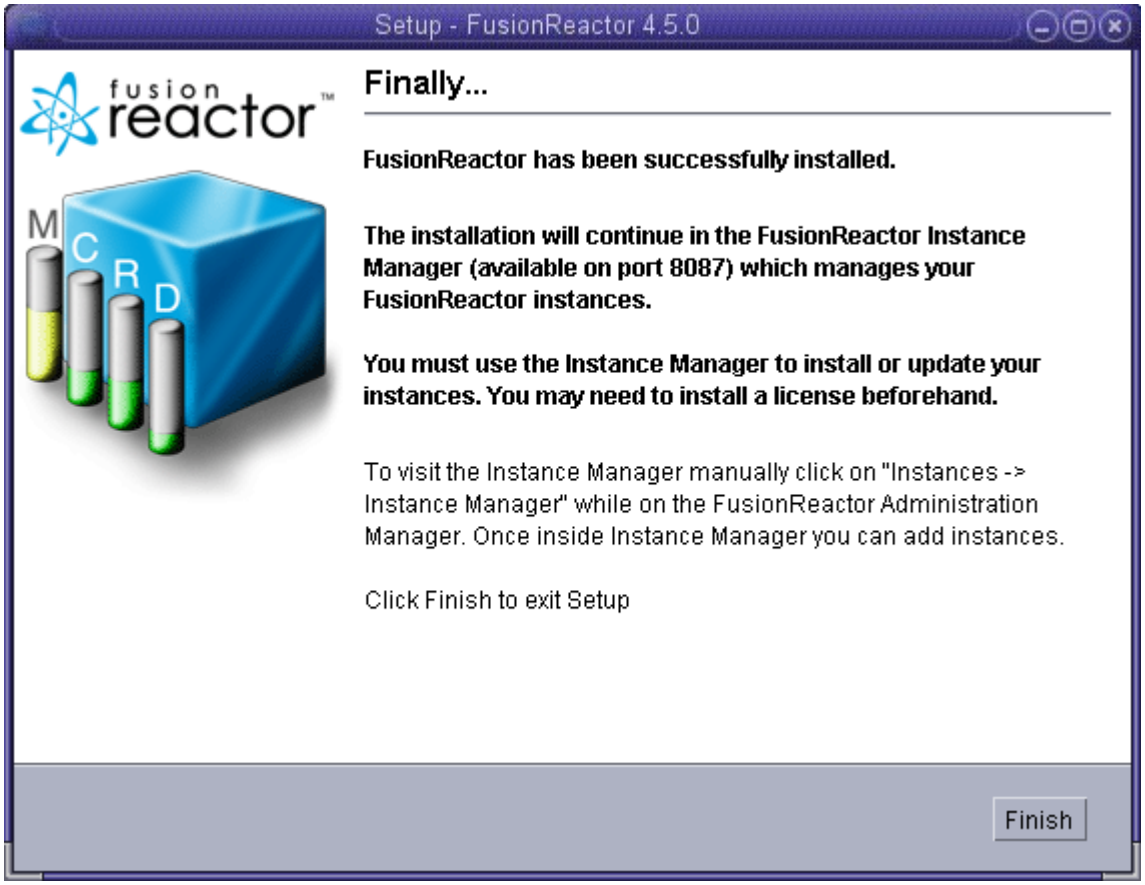


Figure 11: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 11: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 8. This is the last screen in the setup, clicking Finish will close the window.

Updating FusionReactor 3.x on Solaris

Executables

Executable	System Architecture
FusionReactor_solaris_sparc_4_5_0_x64.sh	SPARC 64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

You can run the FusionReactor 4.5.x Setup to update an already installed version of FusionReactor 3.x.

To start the update, login as user *root* and run the appropriate executable as described in the above section **Exe cutables.**

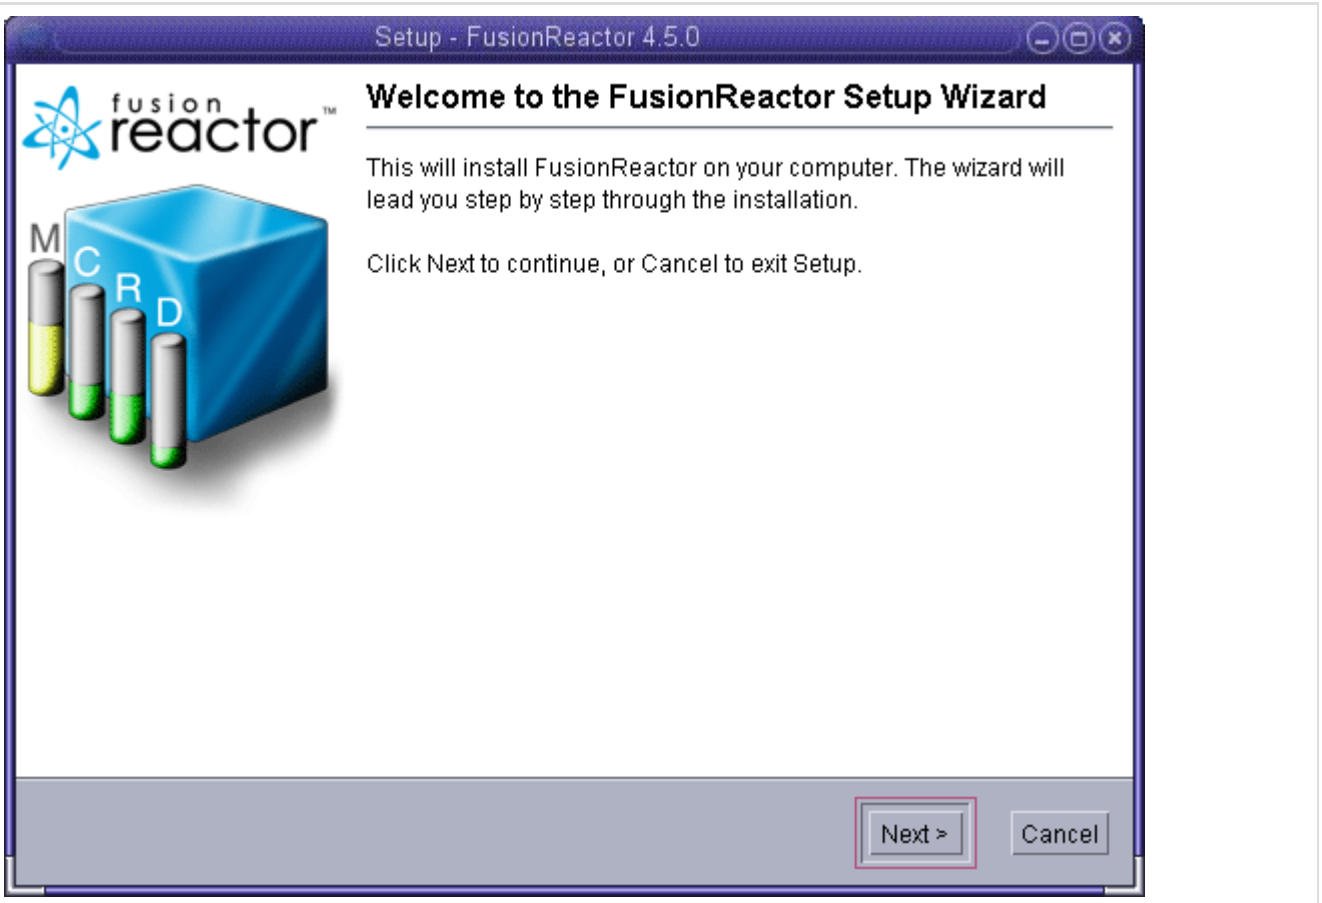


Figure 1: Welcome screen

Figure 1: The first step in the installation wizard is the welcome screen as shown in Figure 1. Clicking the Next button will proceed with the update.

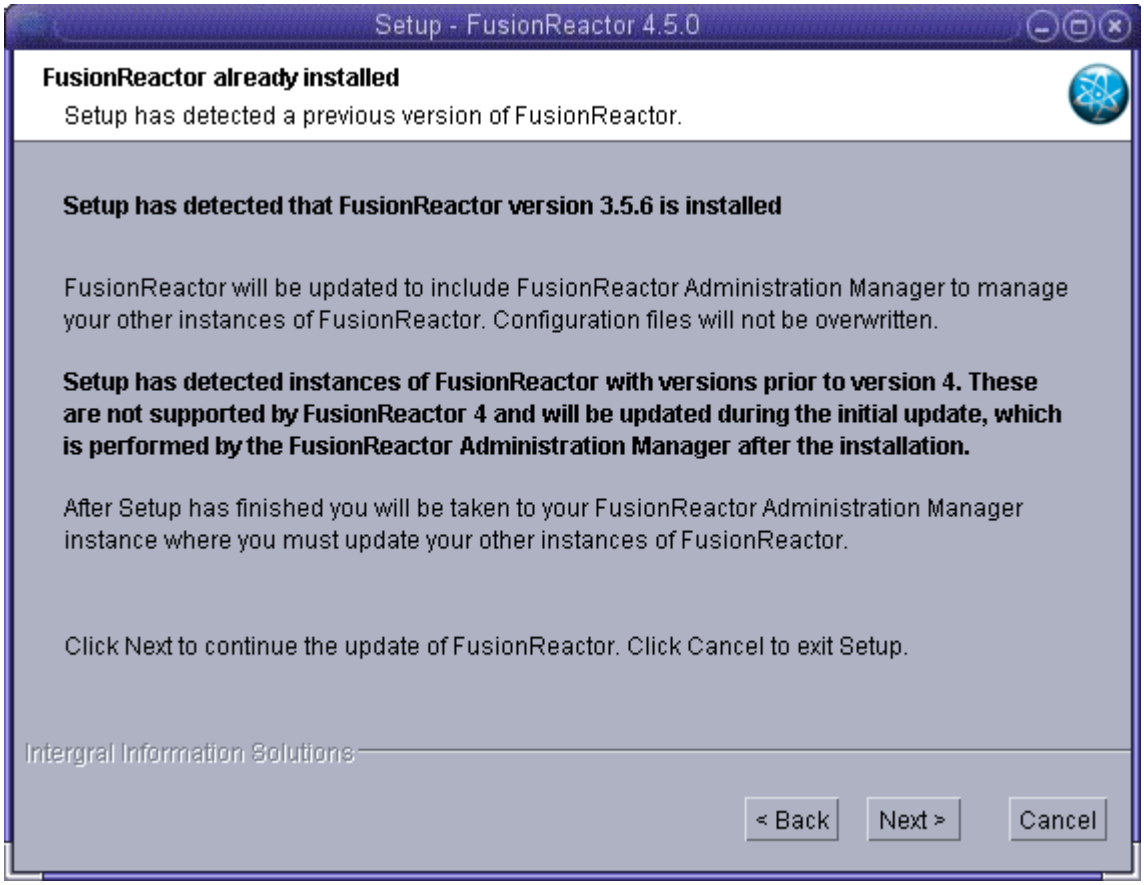


Figure 2: Update the Current Version

Figure 2: If there is already a version of FusionReactor 3.x on your computer that has been installed with the FusionReactor 3.x Setup, you will see a screen similar to the one shown. Carefully read the information and click Next to proceed with the update. Click Cancel to exit the installation wizard.

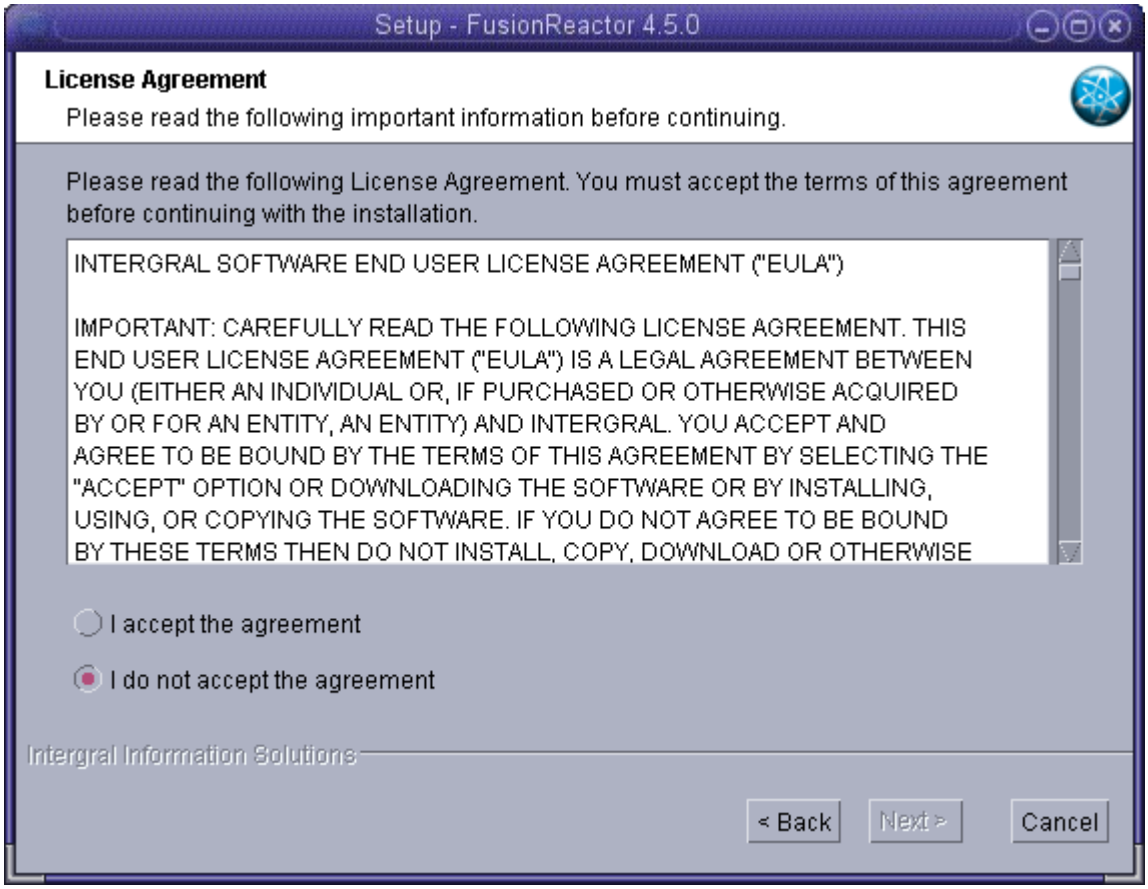


Figure 3: License Agreement

Figure 3: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.

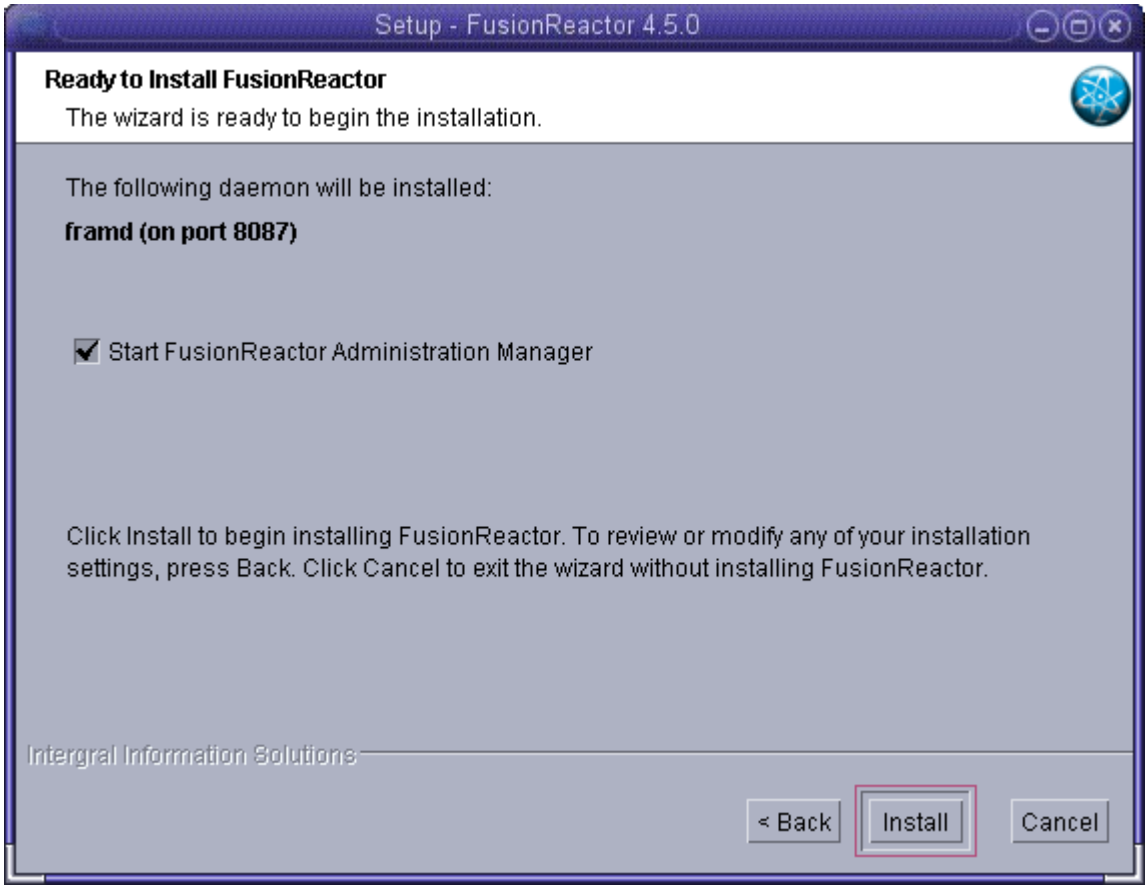


Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the sever ports (See Figure 5). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 6).

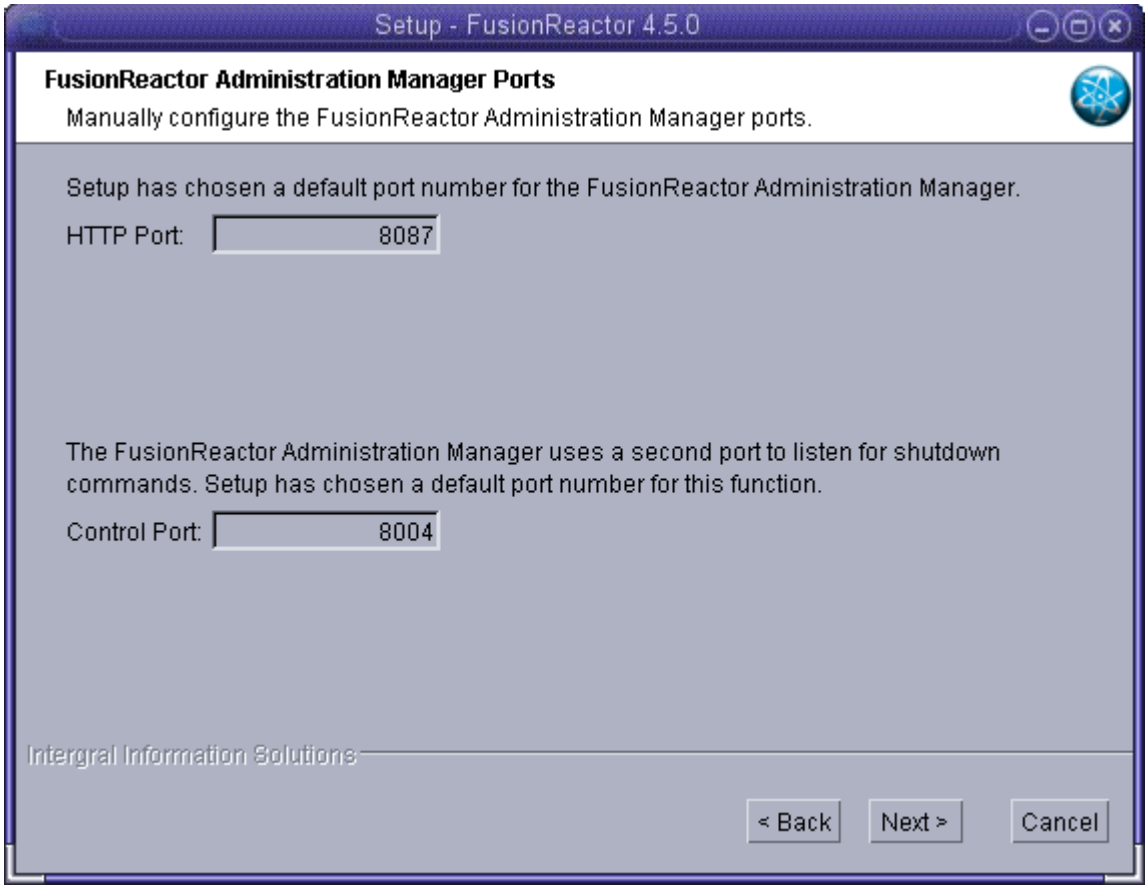


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

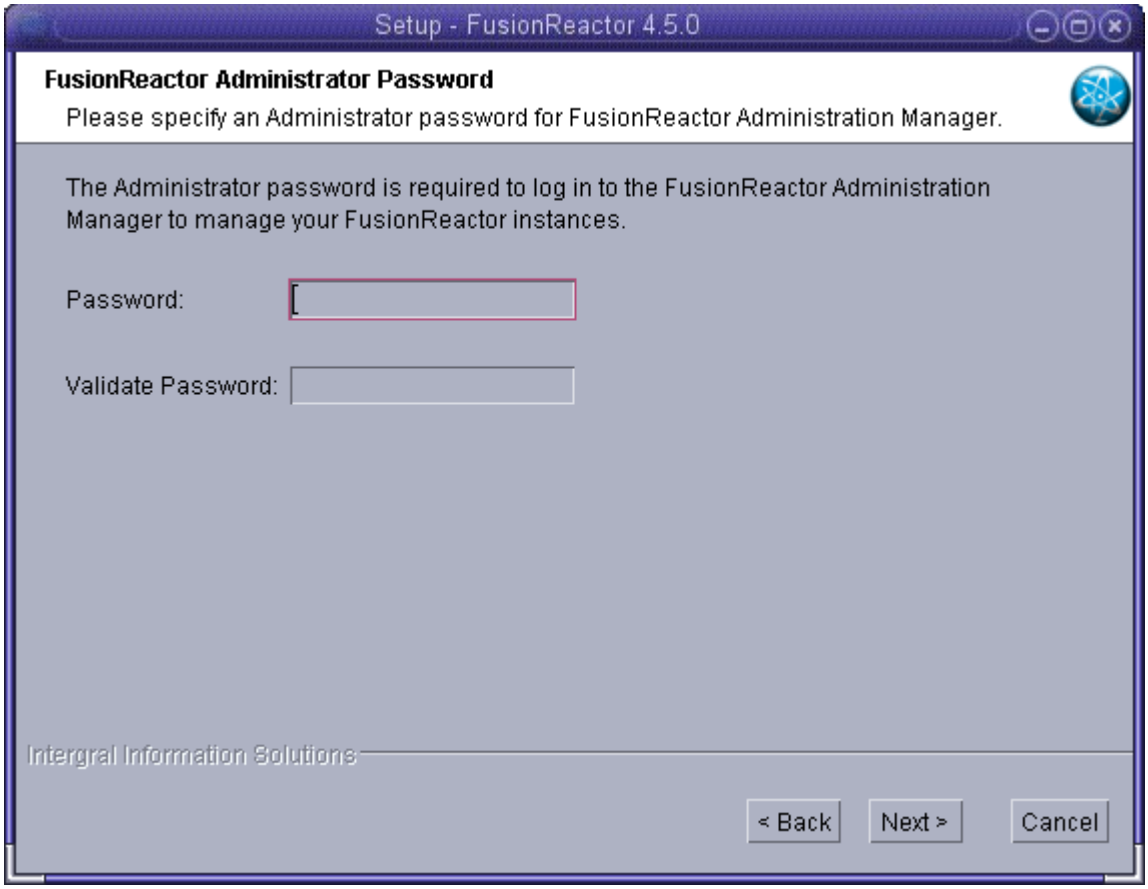


Figure 6: Administration Manager Password

Figure 6: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

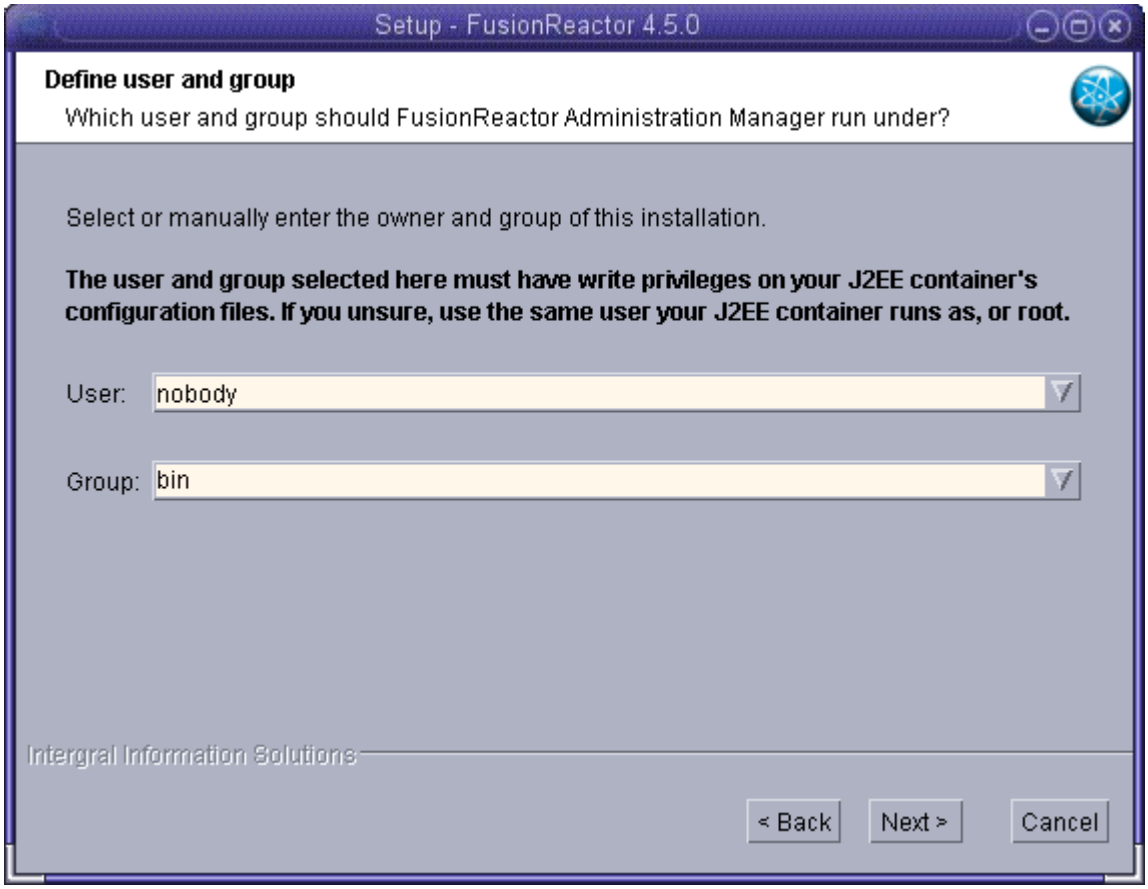


Figure 7: Choose a User and a Group

Figure 7: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM).

Please Note

FusionReactor Administration Manager (FRAM) must be installed as a user who has a real shell and can log in to the system. By default, the installer will pick **nobody/root** for this purpose, but can't presently check that this user can actually log in. If your **nobody** user does not have a shell, please change their shell to **/bin/sh** (or your preferred shell) or select another user.

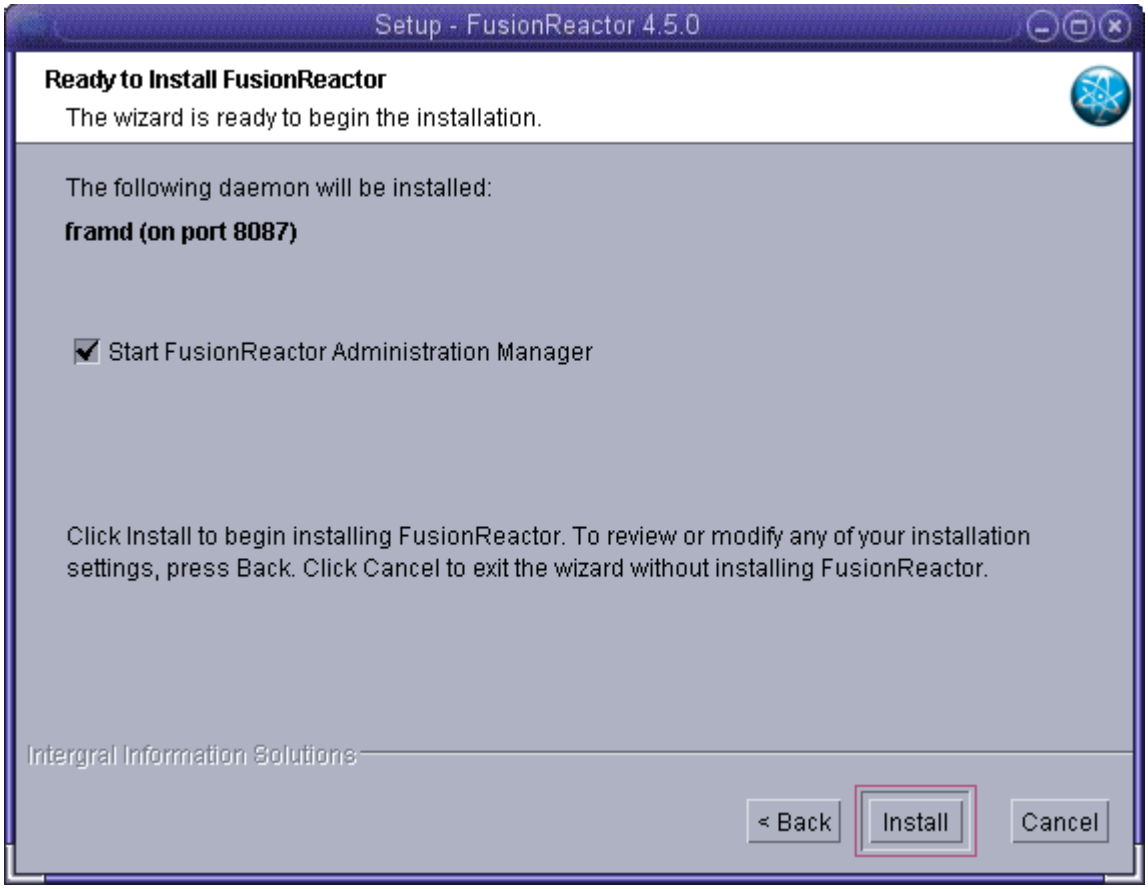


Figure 8: Ready to Install

Figure 8: Setup can automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin updating to FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

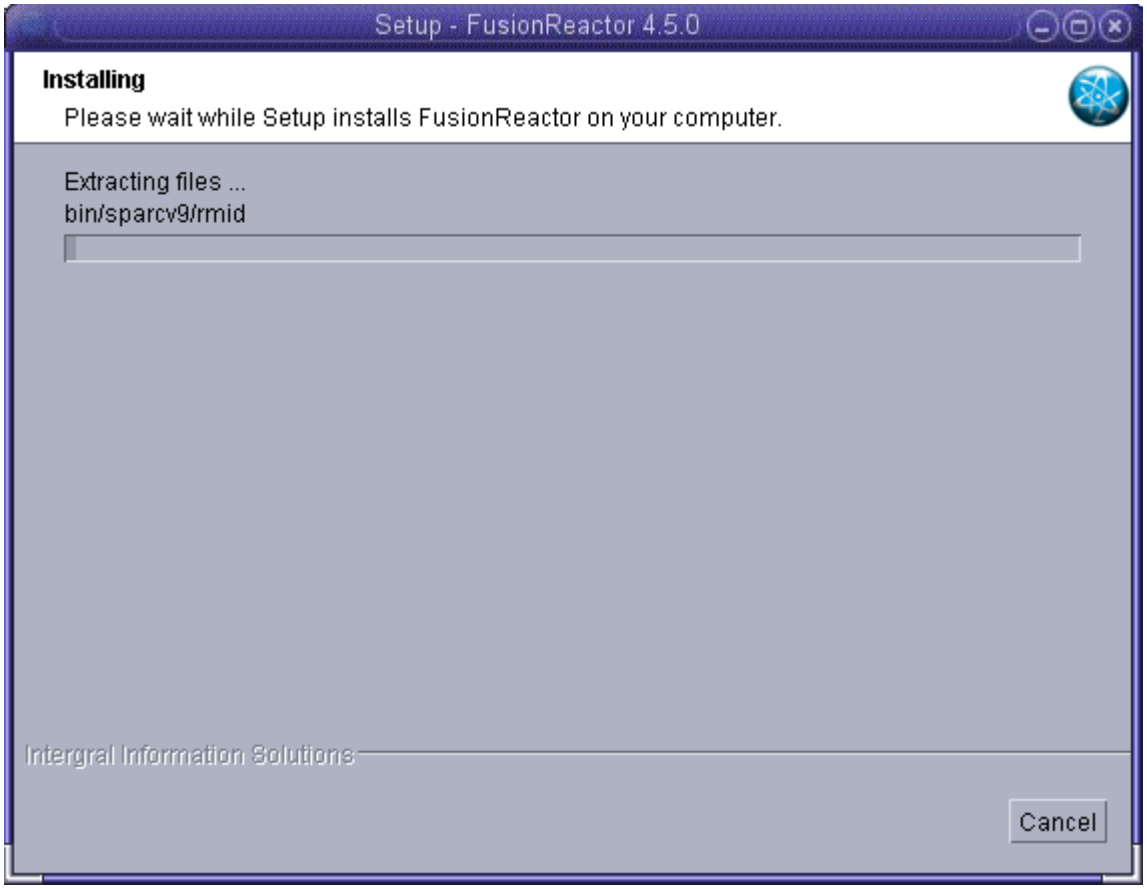


Figure 9: Installation Progress

Figure 9: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 10: Configure Instances Screen

Figure 10: After the installation has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. To configure now launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Server Scan Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close"* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 8.

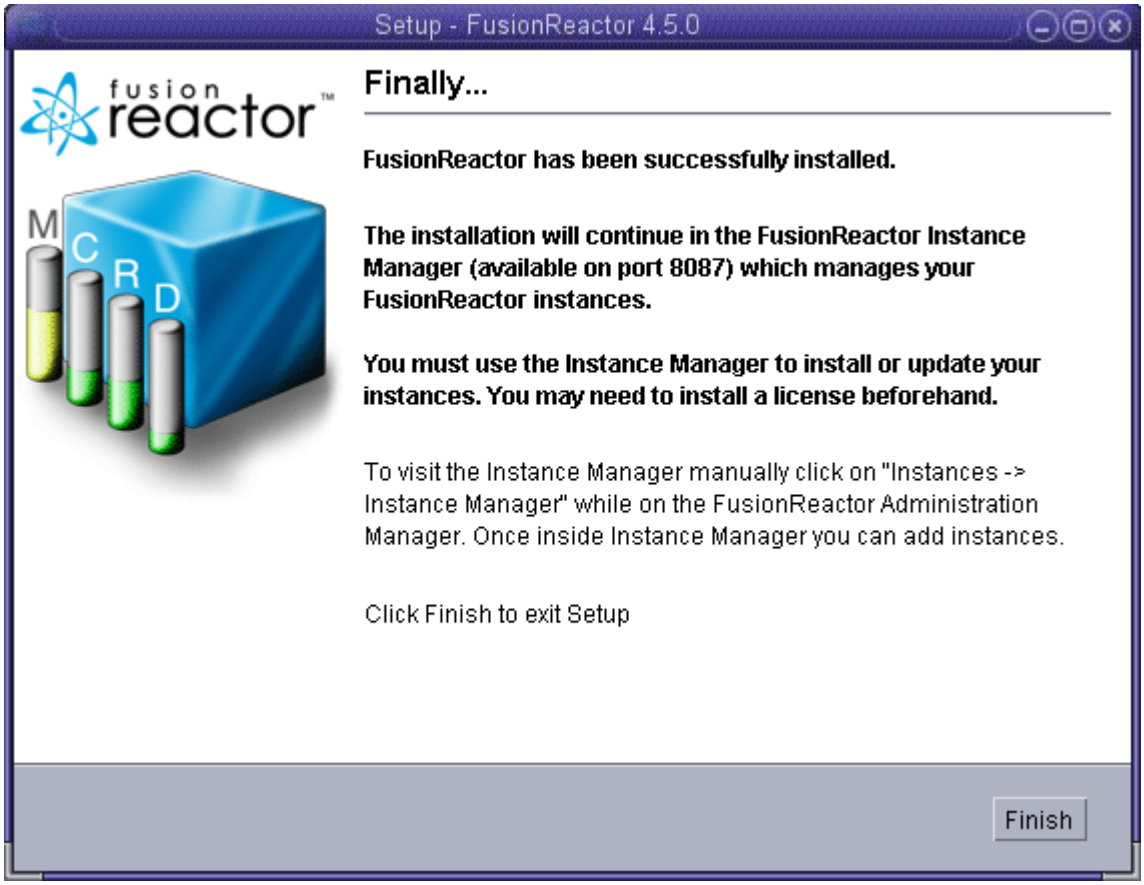


Figure 11: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 11: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 8. This is the last screen in the setup, clicking Finish will close the window.

Updating FusionReactor 2.0 on Solaris

Upgrade Not Possible

FusionReactor 4.x.x does not support upgrades from FusionReactor 2.x.

If you currently have FusionReactor 2.x installed it is not possible to upgrade to FusionReactor 4.x.x. When running the FusionReactor 4.x.x setup you will be presented with a screen similar to that shown in Figure 1 below.

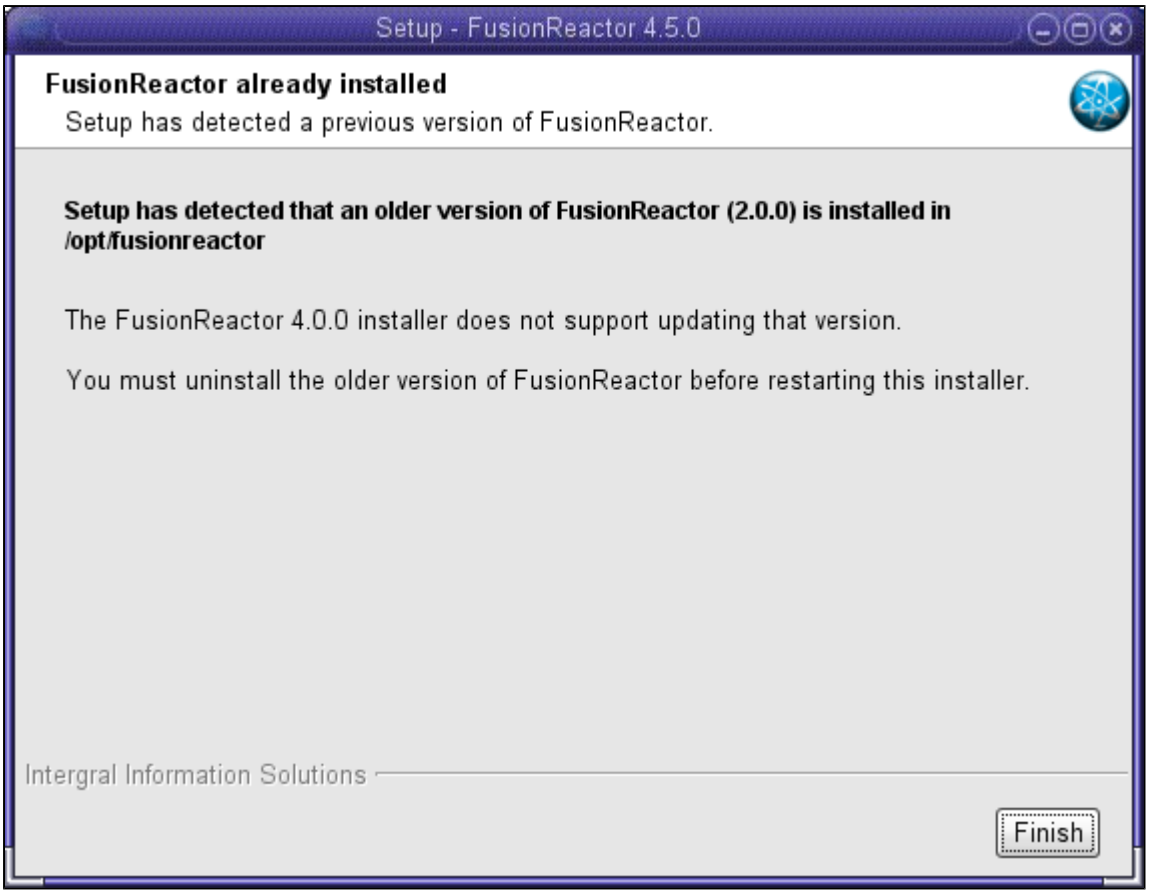


Figure 1: Update Not Supported

To install FusionReactor 4.5.x you must first uninstall FusionReactor 2.x and then run the FusionReactor 4.5.x Setup again. Please follow the guide on how to Uninstall FusionReactor 2.x on Solaris

Updating FusionReactor 1.0 on Solaris

> **Upgrade Not Possible**
>
> FusionReactor 4.x.x does not support upgrading from FusionReactor 1.x.

If you currently have FusionReactor 1.x installed it is not possible to upgrade to FusionReactor 4.x.x.

To install FusionReactor 4.5.x you must first uninstall FusionReactor 1.x and then run the FusionReactor 4.5.x Setup again.

Console Installation on Solaris

Executables

Executable	System Architecture
FusionReactor_solaris_sparc_4_5_0_x64.sh	SPARC 64-bit

The FusionReactor setup comes bundled with an appropriate Java Runtime Environment (JRE) so make sure to select the correct executable.



Setup

The setup needs to be run as root.

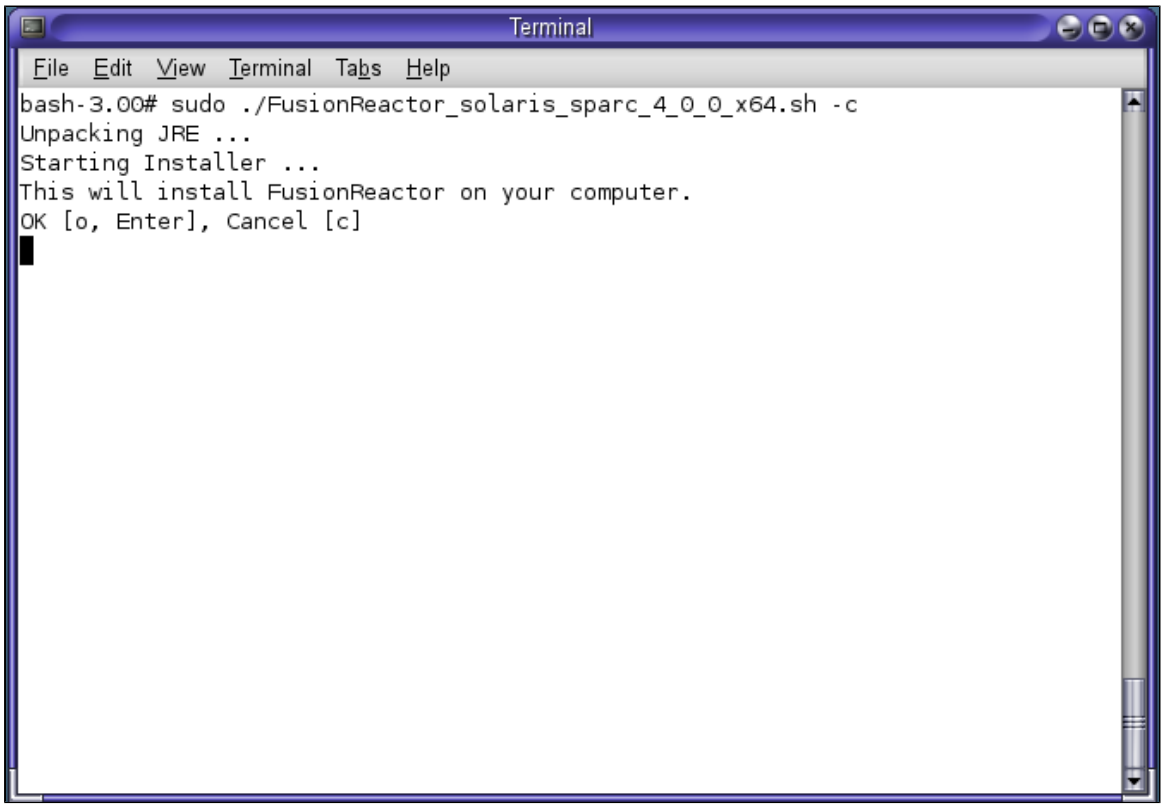


Figure 1: Start the Installation

Figure 1: To run the FusionReactor 4.5.0 installation in console mode, you need to add the c switch .
For example: `sudo ./FusionReactor_linux_4_5_0.sh -c`

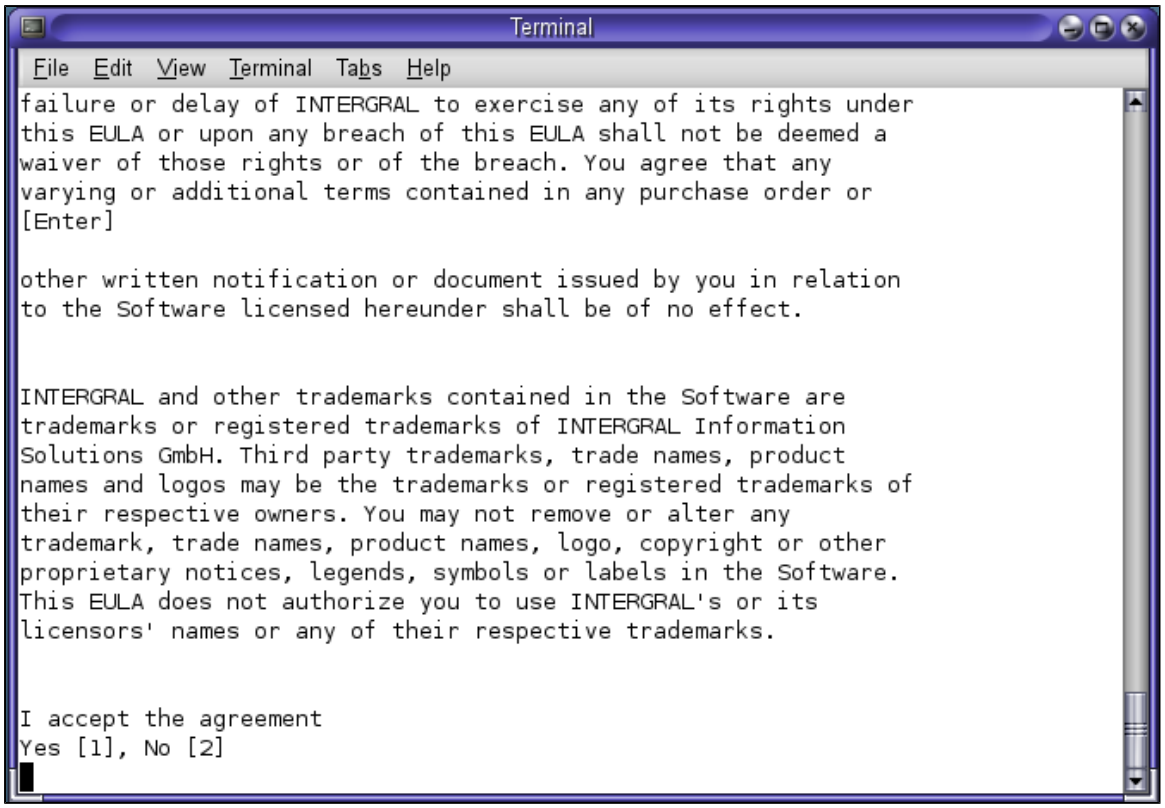


Figure 2: License Agreement

Figure 2: After starting the install you will be shown the license agreement. Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), by entering '1'.

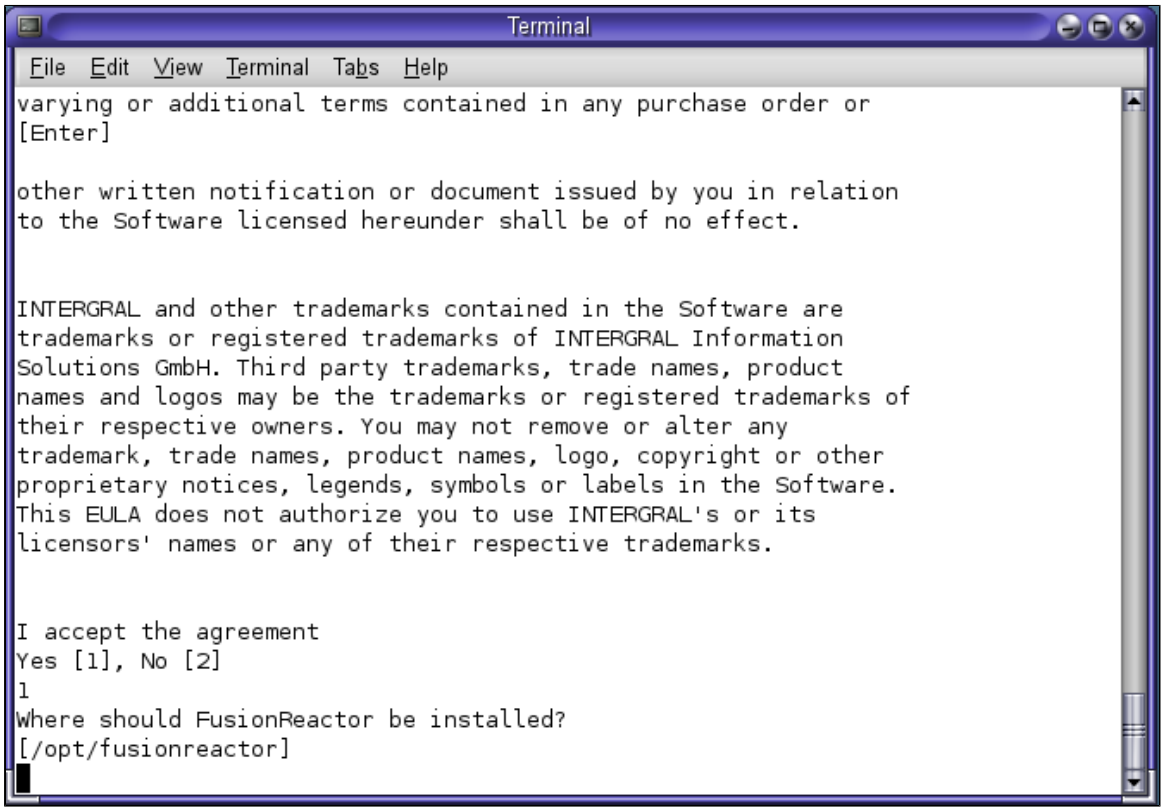


Figure 3: Installation Directory Selection

Figure 3: Now you have to select where FusionReactor is to be installed. By default (pressing enter directly),

FusionReactor will be installed into `/opt/fusionreactor/`.

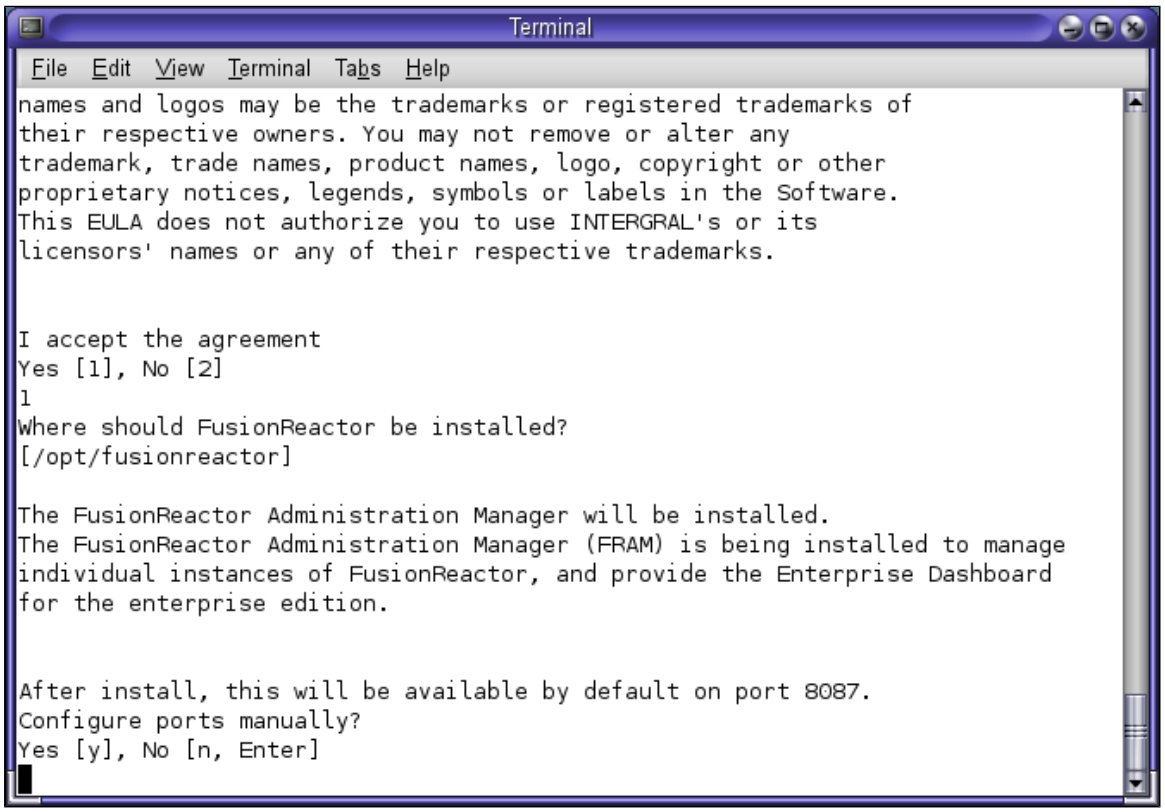


Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: You will now be asked if you want to configure the ports manually for FusionReactor Administration Manager (FRAM). Either enter 'y' to configure them (see Figure 5 onwards) or enter 'n' or just Enter on the keyboard (see Figure 7 onwards).

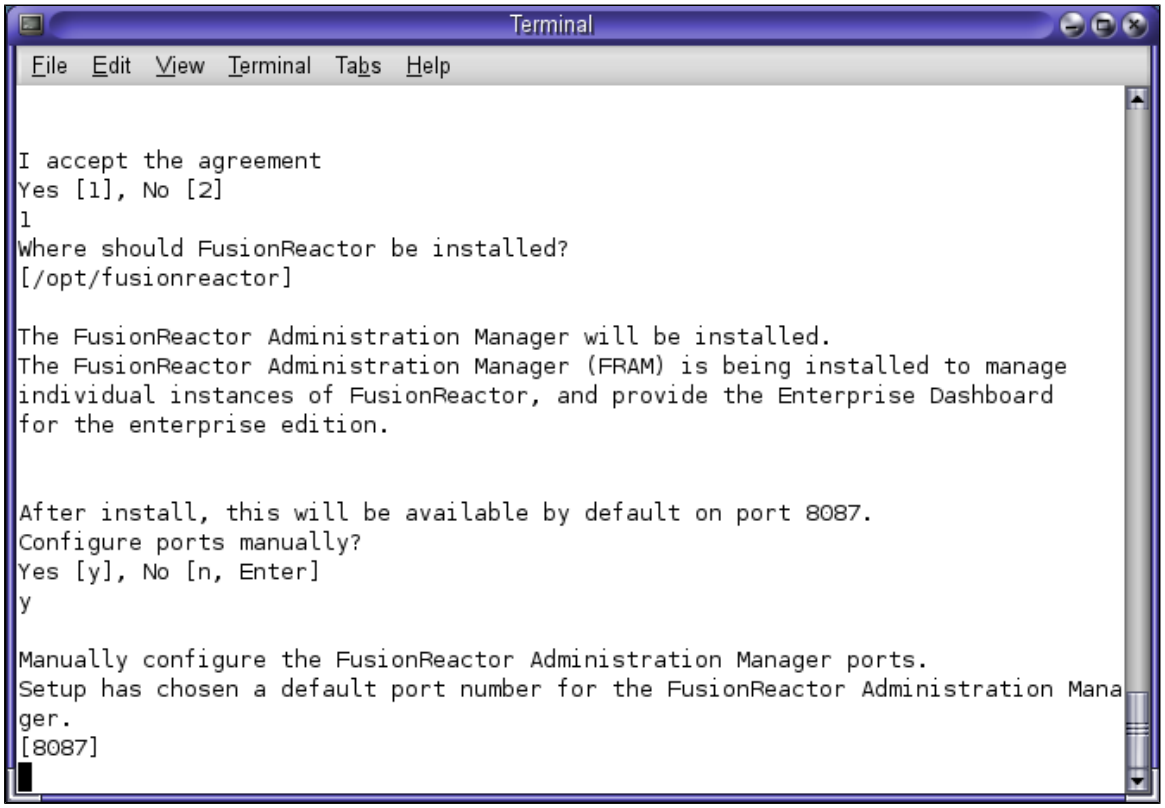


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (the port you use to view FusionReactor in a web browser). It is important not to use ports that are already in use by other applications. NOTE that the default port for FRAM is 8087.

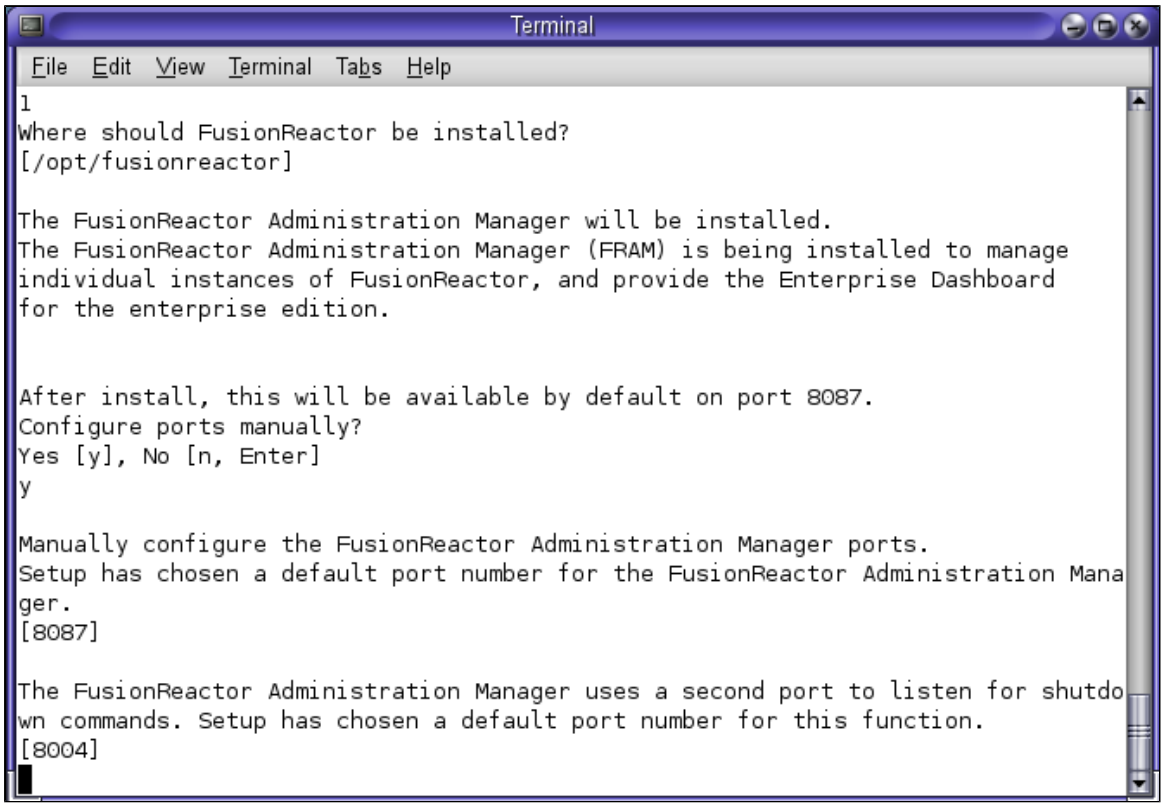


Figure 6: Manually Configuring Ports

Figure 6: You must also configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications.

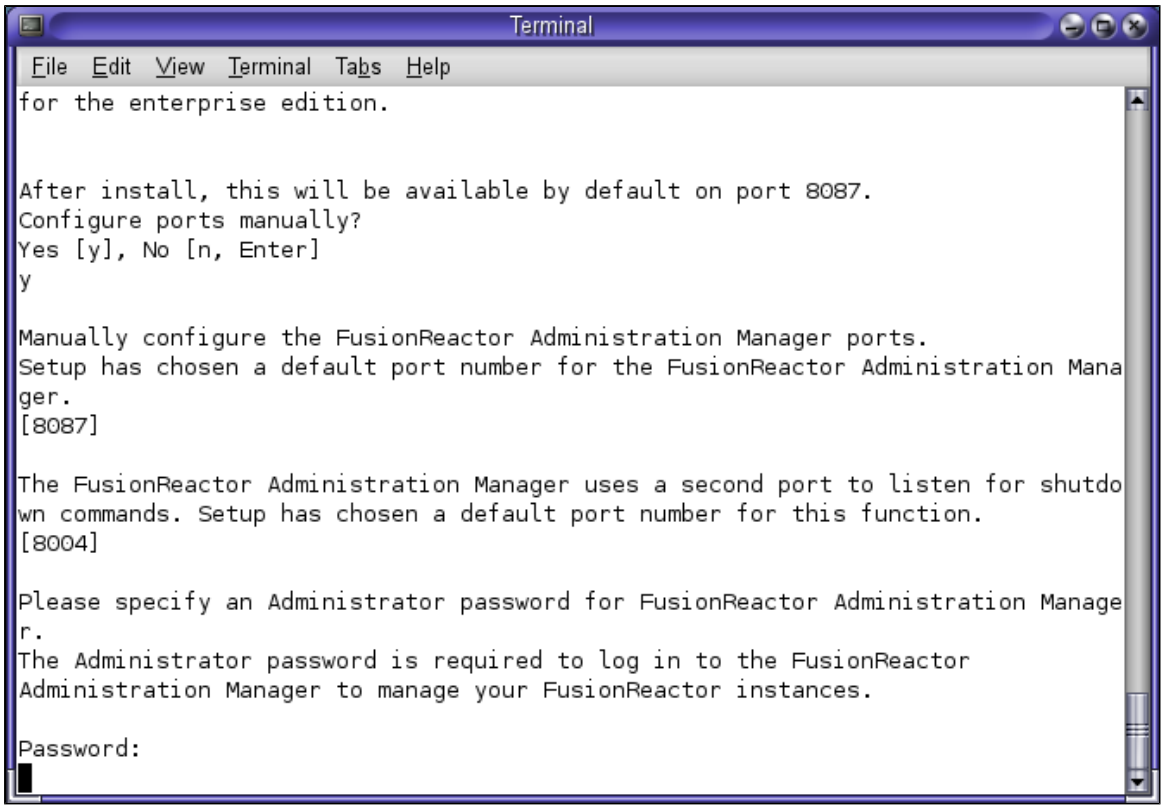


Figure 7: Administration Manager Password

Figure 7: Here you will be asked to specify a password for FRAM's Administrator account. It must be longer than 5 characters.

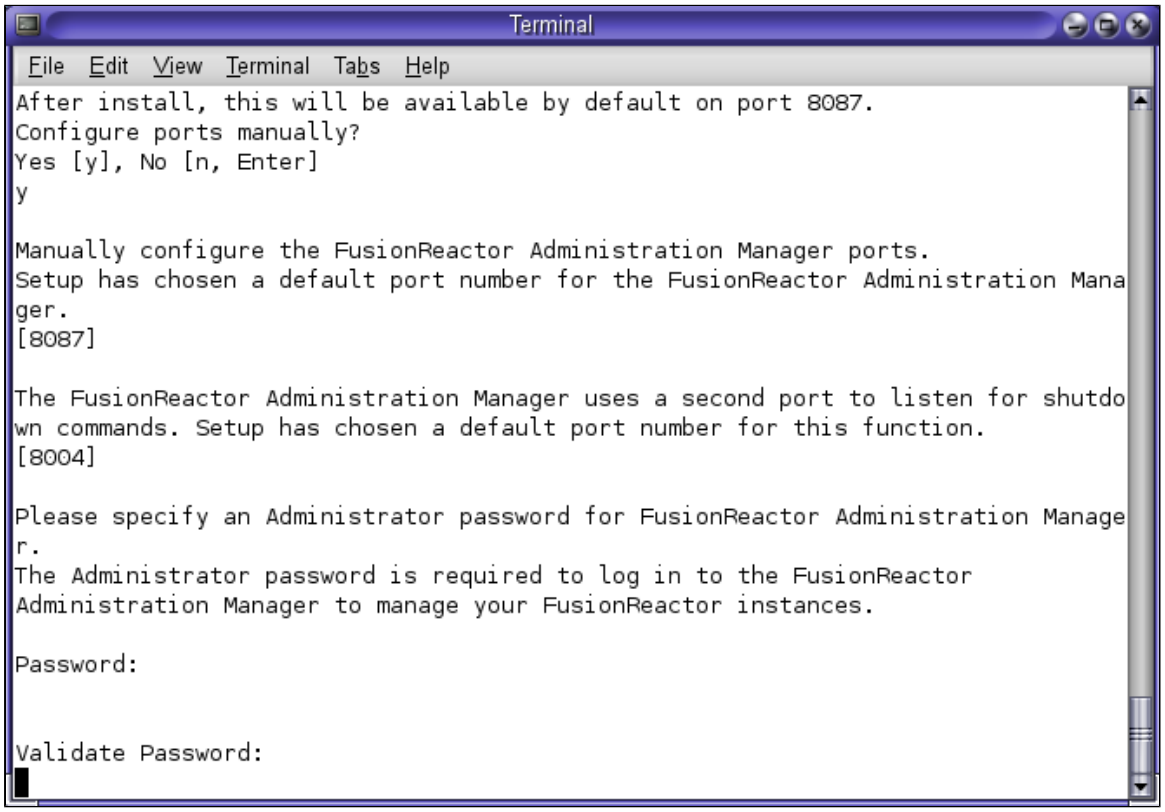


Figure 8: Administration Manager Password

Figure 8: You will be asked to validate your password.

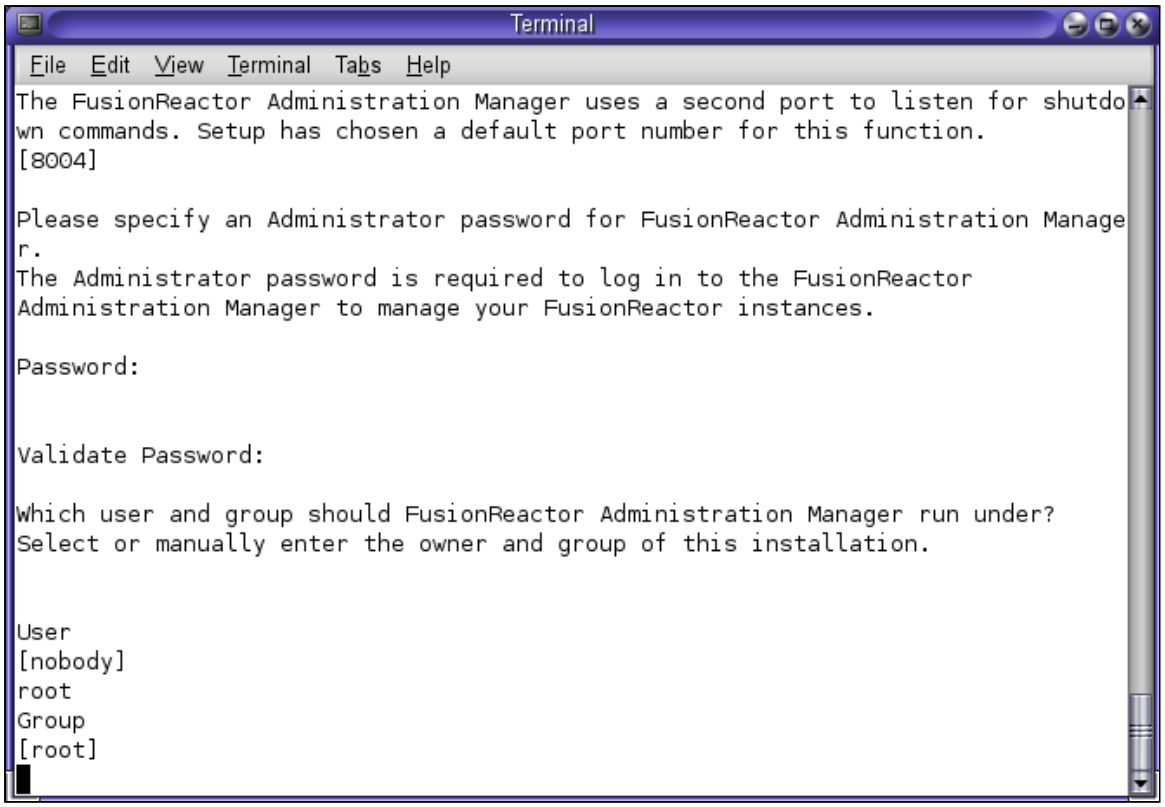


Figure 9: Choose a User and a Group

Figure 9: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM). Only that user or someone with the same or higher privileges will be allowed to run FRAM.

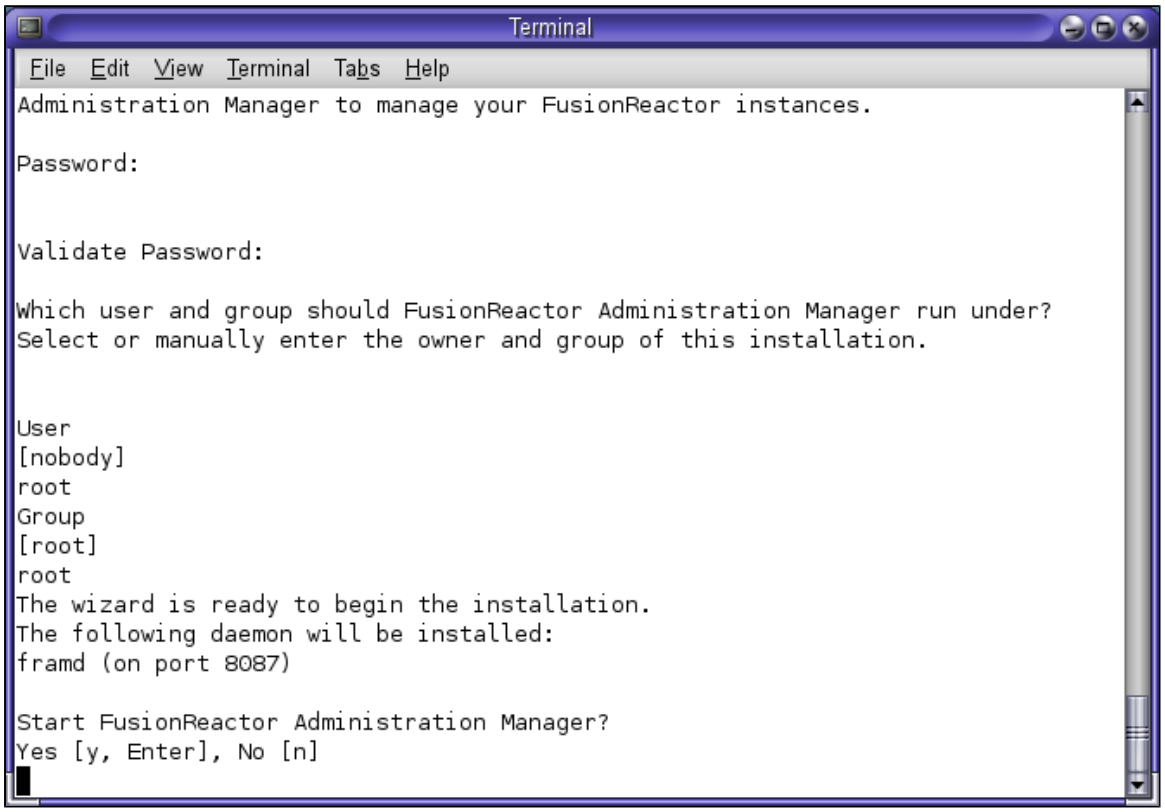


Figure 10: Ready to Install

Figure 10: Here you will be given the choice of starting up the FRAM daemon. Either entering 'y' or pressing enter will start the daemon while entering 'n' will not start the daemon. FRAM can be started manually after the

installation has finished.

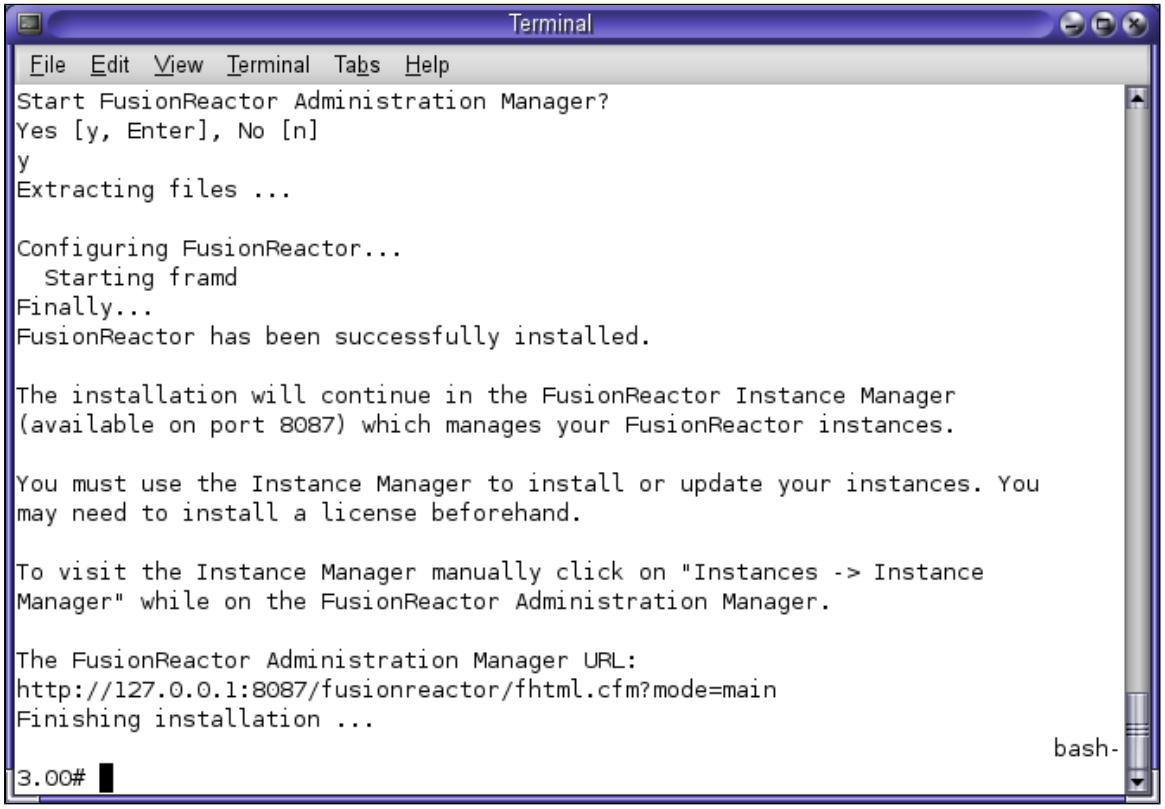


Figure 11: Finish Screen

Figure 11: You have successfully installed FusionReactor. Open the URL to go to FRAM.

Installing FusionReactor on Mac

Installing FusionReactor 4.5.x on Mac

Updating FusionReactor 3.x on Mac

Installing FusionReactor 4.5.x on Mac

Executables

Executable	System Architecture
FusionReactor_macos_4_5_0.dmg	32-bit
FusionReactor_macos_4_5_0.dmg	64-bit

Setup



Figure 1: Welcome screen

Figure 1: The first step in the installation wizard is the welcome screen.

Clicking the Next button will proceed with the installation and open the License Agreement as shown in Figure 2. Clicking Cancel will exit the Setup.

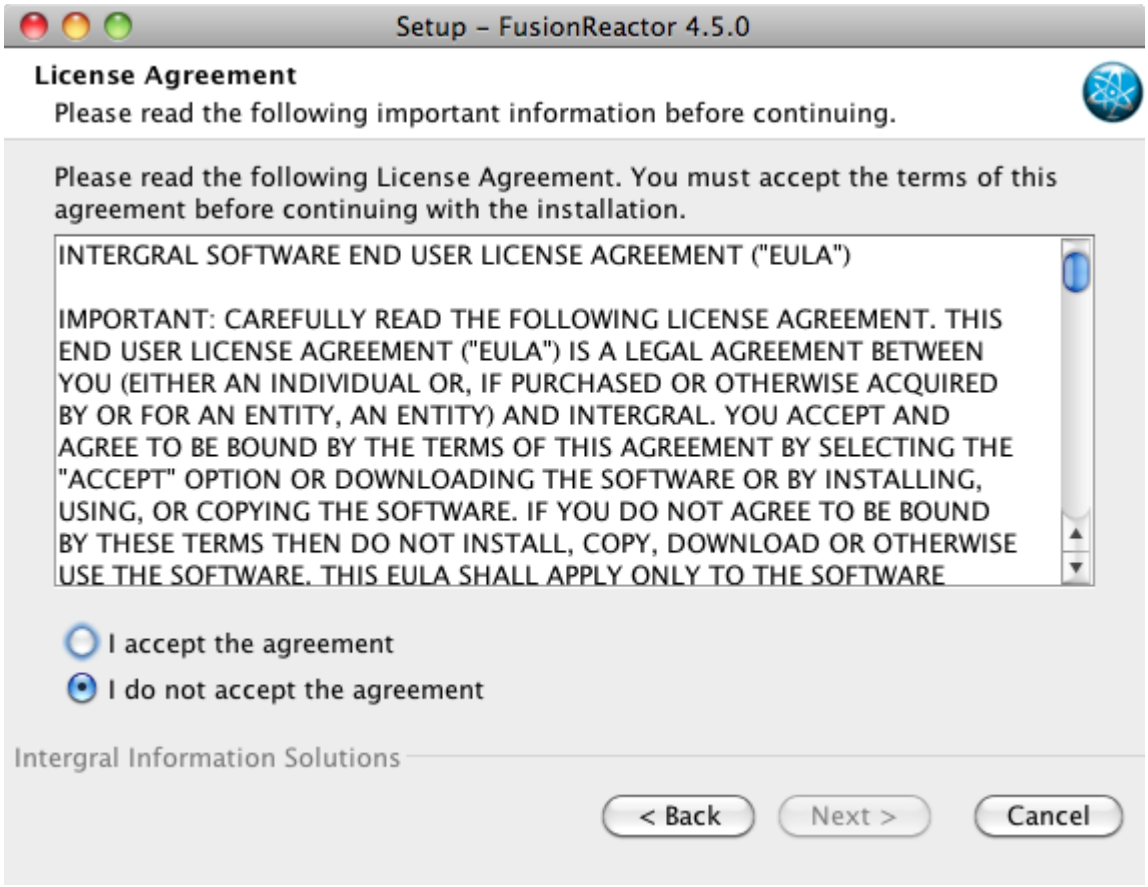


Figure 2: License Agreement

Figure 2: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.

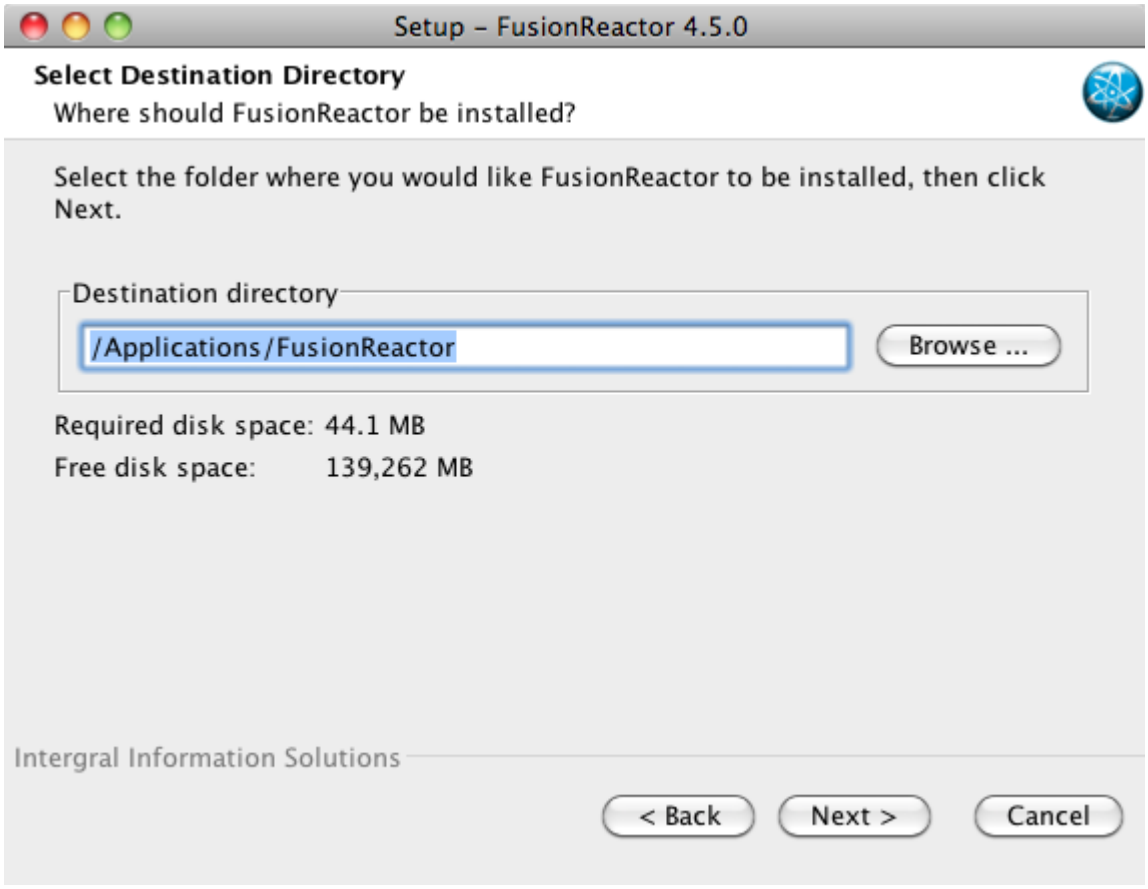


Figure 3: Installation Directory Selection

Figure 3: Now you have to select the directory in which FusionReactor should be installed. By default FusionReactor will be installed into `/Applications/FusionReactor/`



Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the server ports (See Figure 5). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 6).

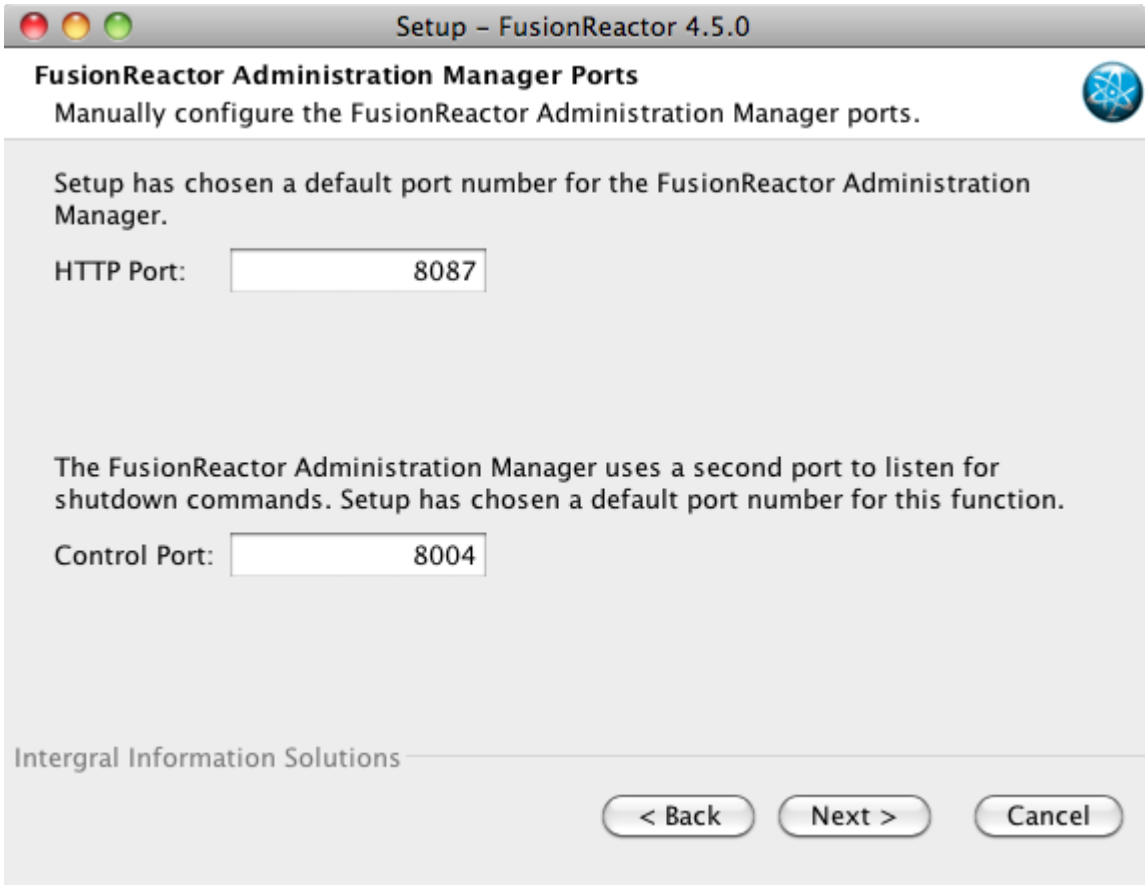


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

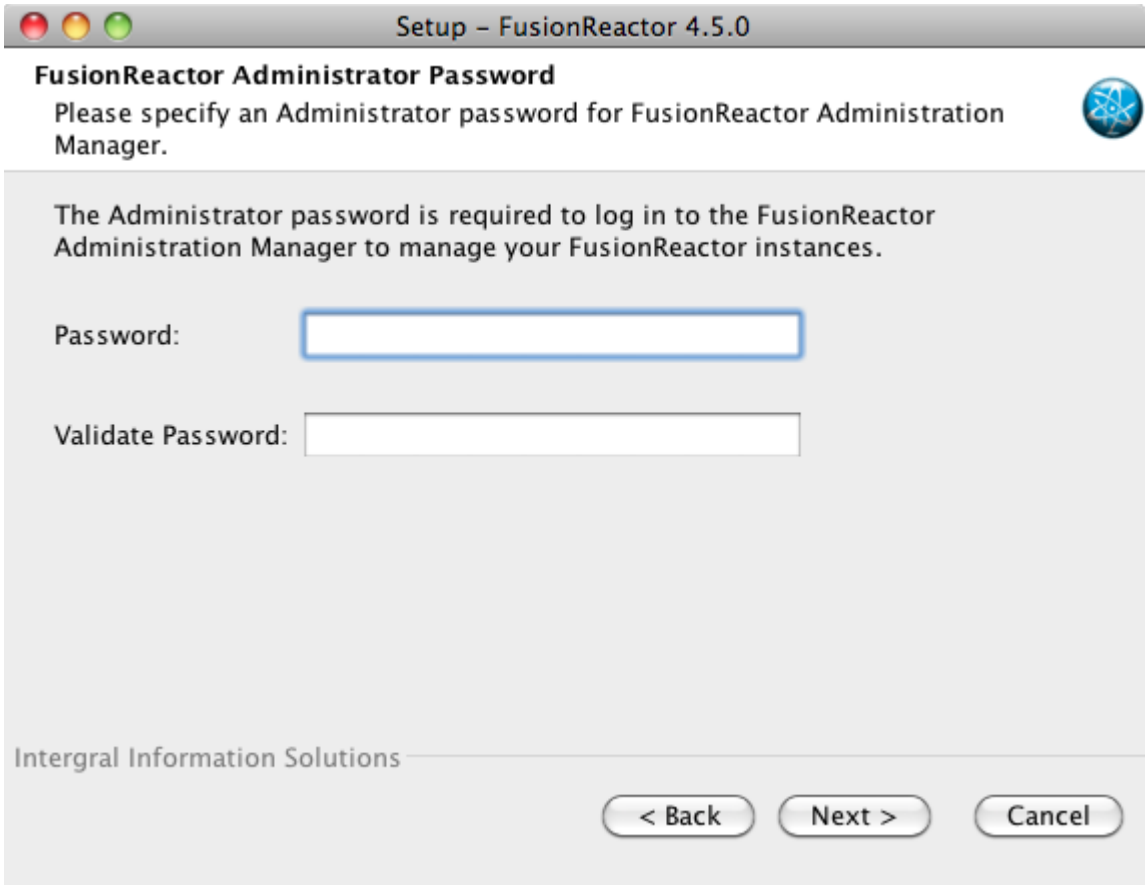


Figure 6: Administration Manager Password

Figure 6: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

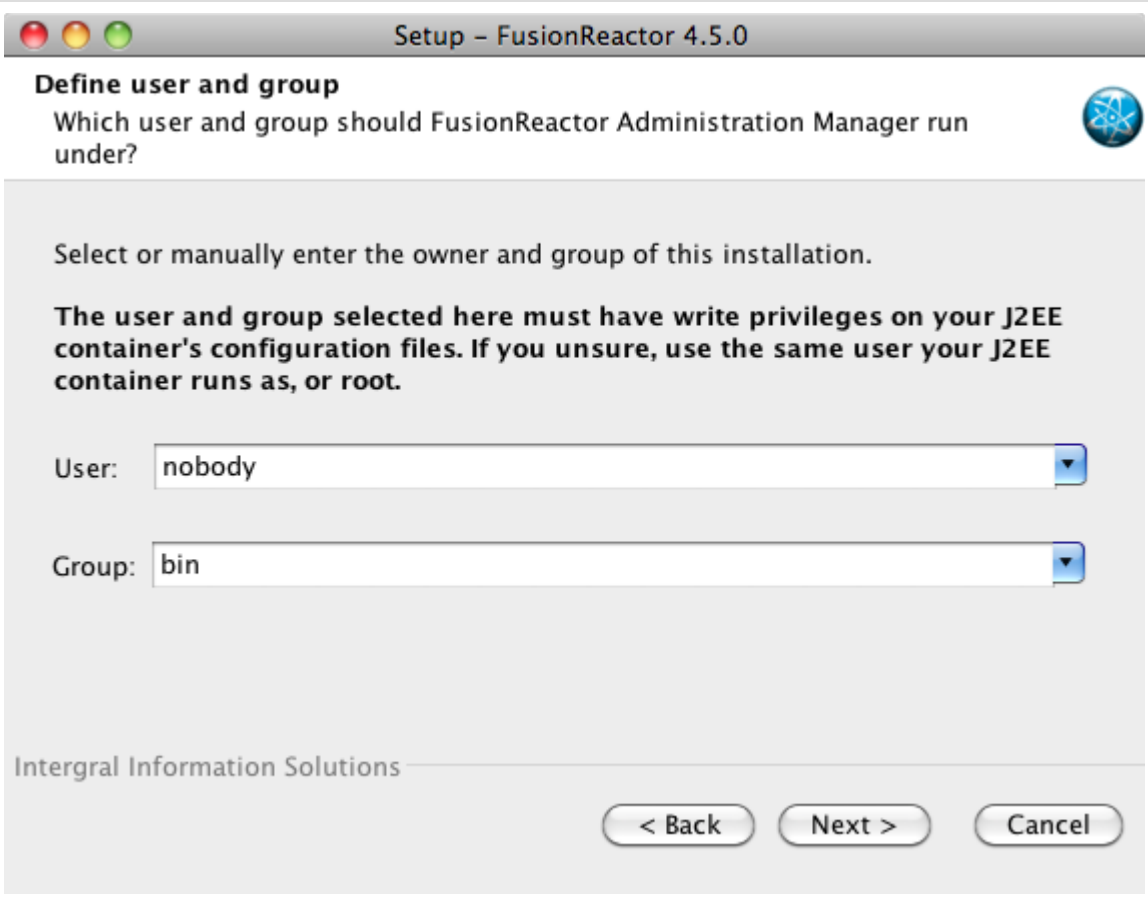


Figure 7: Choose a User and a Group

Figure 7: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM). Only that user or someone with the same or higher privileges will be allowed to run FRAM.

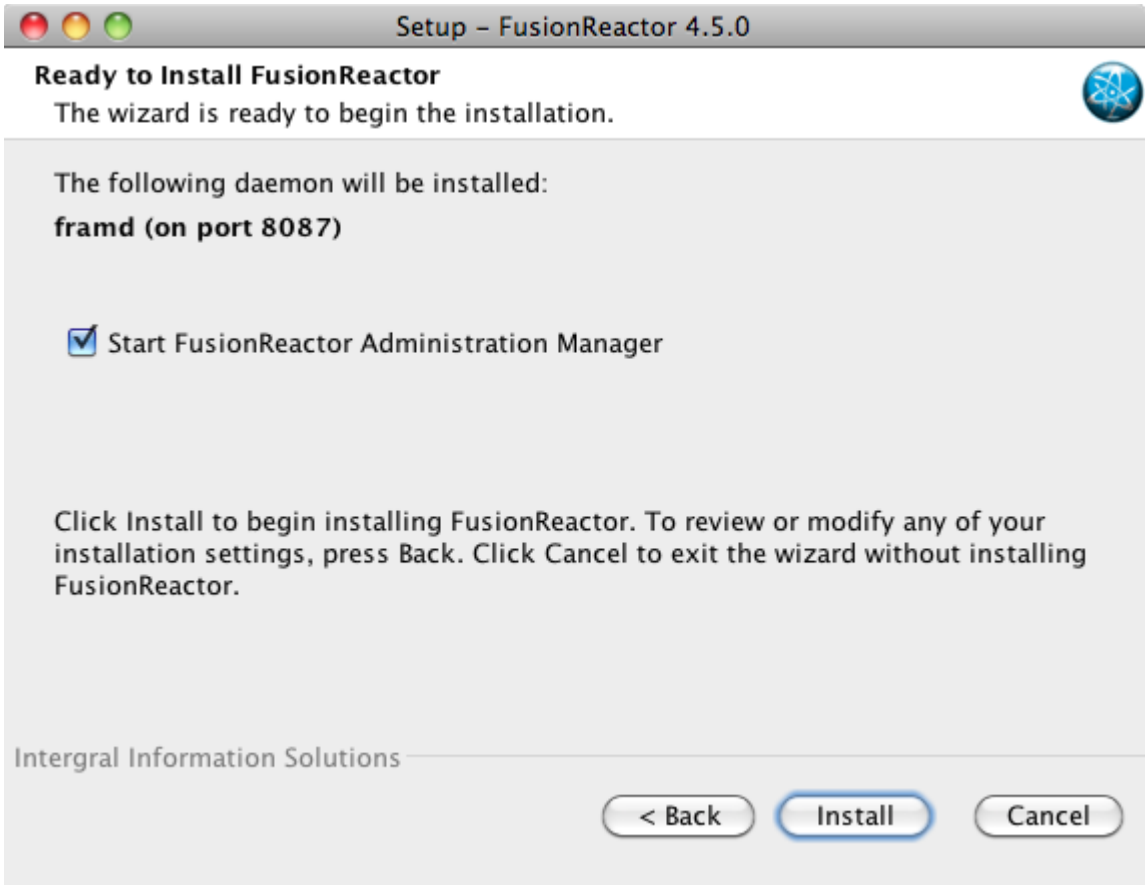


Figure 8: Ready to Install

Figure 8: Setup can automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin installing FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

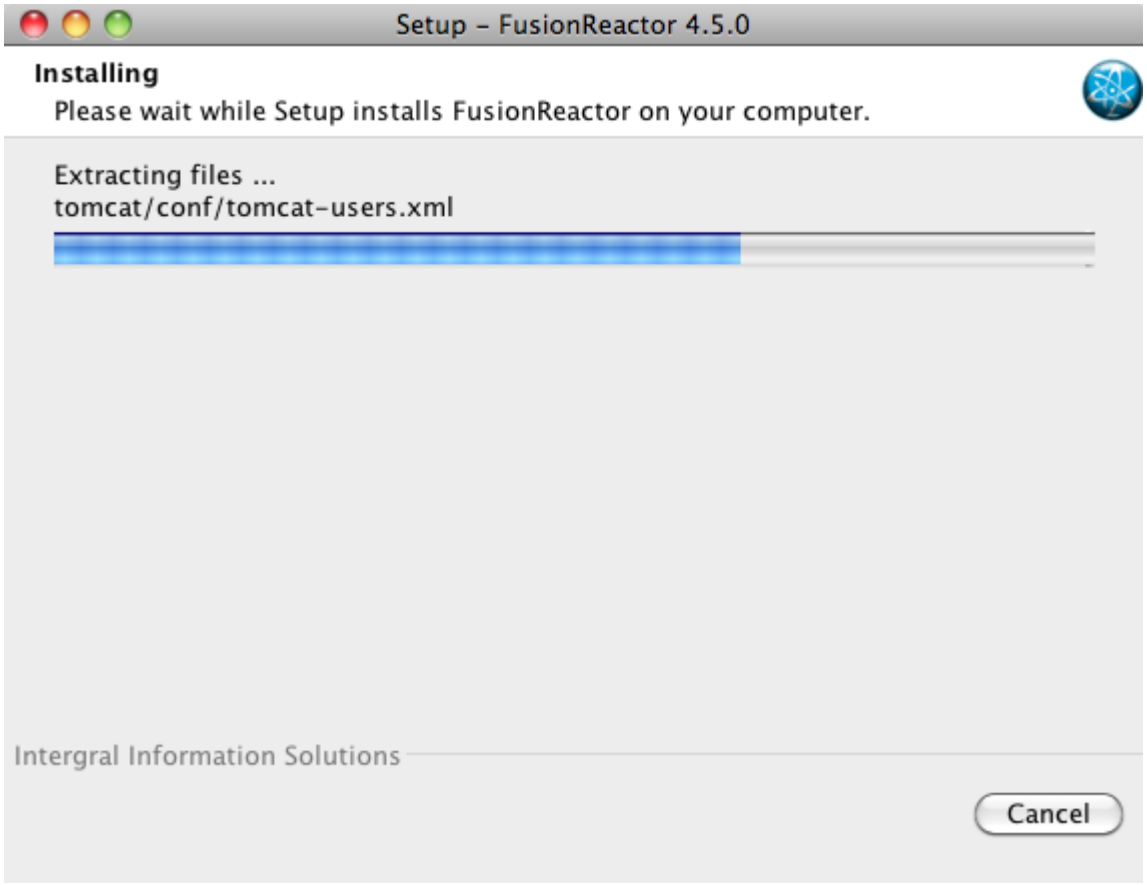


Figure 9: Installation Progress

Figure 9: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 10: Configure Instances Screen

Figure 10: After the installation has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. To configure now launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Server Scan Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close"* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 8.



Figure 11: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 11: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 8. This is the last screen in the setup, clicking Finish will close the window.

Updating FusionReactor 3.x on Mac

Executables

Executable	System Architecture
FusionReactor_macos_4_5_0.dmg	32-bit
FusionReactor_macos_4_5_0.dmg	64-bit

Setup

You can run the FusionReactor 4.5.x Setup to update an already installed version of FusionReactor 3.x.



Figure 1: Welcome screen

Figure 1: The first step in the installation wizard is the welcome screen as shown in Figure 1. Clicking the Next button will proceed with the update.

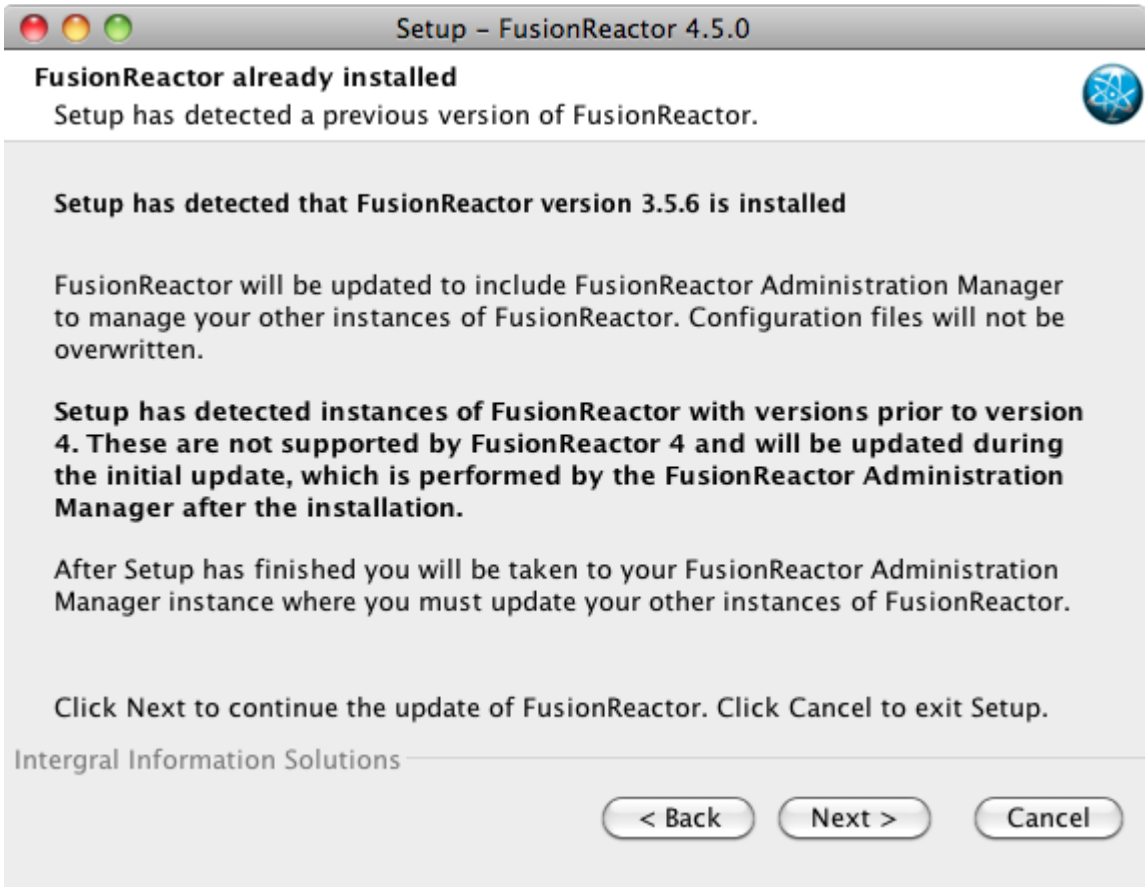


Figure 2: Update the Current Version

Figure 2: If there is already a version of FusionReactor 3.x on your computer that has been installed with the FusionReactor 3.x Setup, you will see a screen similar to the one shown. Carefully read the information and click Next to proceed with the update. Click Cancel to exit the installation wizard.

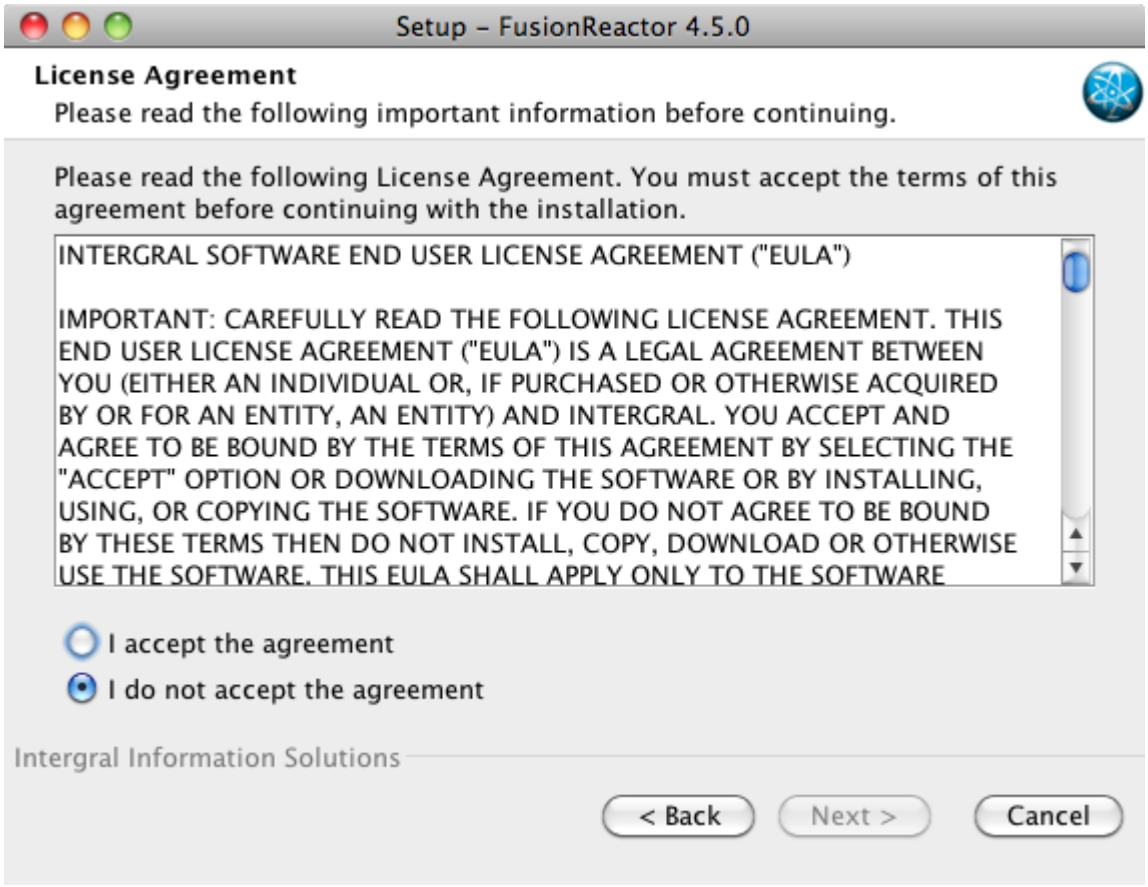


Figure 3: License Agreement

Figure 3: Please carefully review the license agreement, then choose the '*I accept the agreement*' option (if you accept the terms and conditions), click Next to proceed. Press Cancel if you do not accept the license agreement to exit the setup wizard.



Figure 4: FusionReactor Administration Manager (FRAM)

Figure 4: The default HTTP port for FRAM is 8087. If you wish to manually specify the ports used by FRAM select the *'Configure ports manually'* option, then click Next to configure the server ports (See Figure 5). If you wish to use the default port settings leave the box unchecked and click Next (See Figure 6).

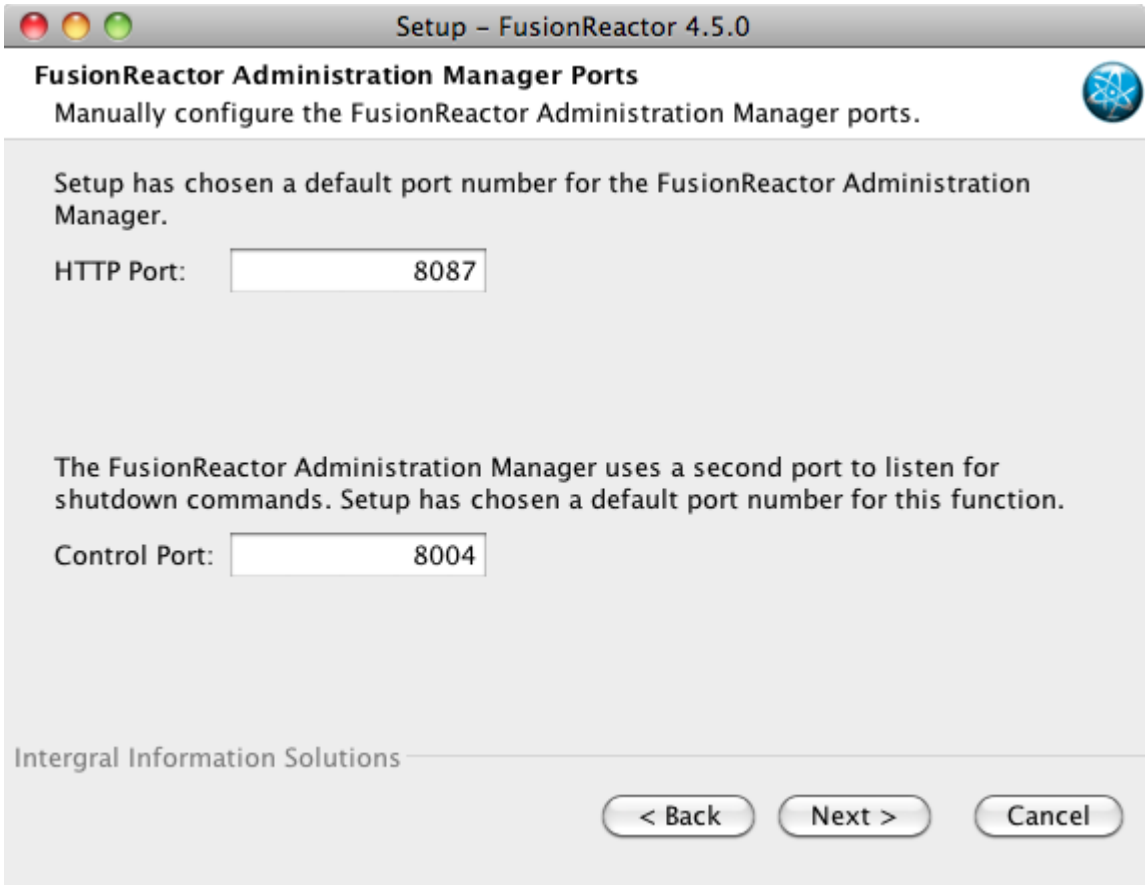


Figure 5: Manually Configuring Ports

Figure 5: Here you can configure the HTTP Port (port to view FusionReactor in a web browser) as well as configure the Control Port (used to listen for shutdown commands). It is important not to use ports that are already in use by other applications. Click Next to use the port configuration you specified. If you wish to use the default port configuration click Back and uncheck the *'Configure ports manually'* option.

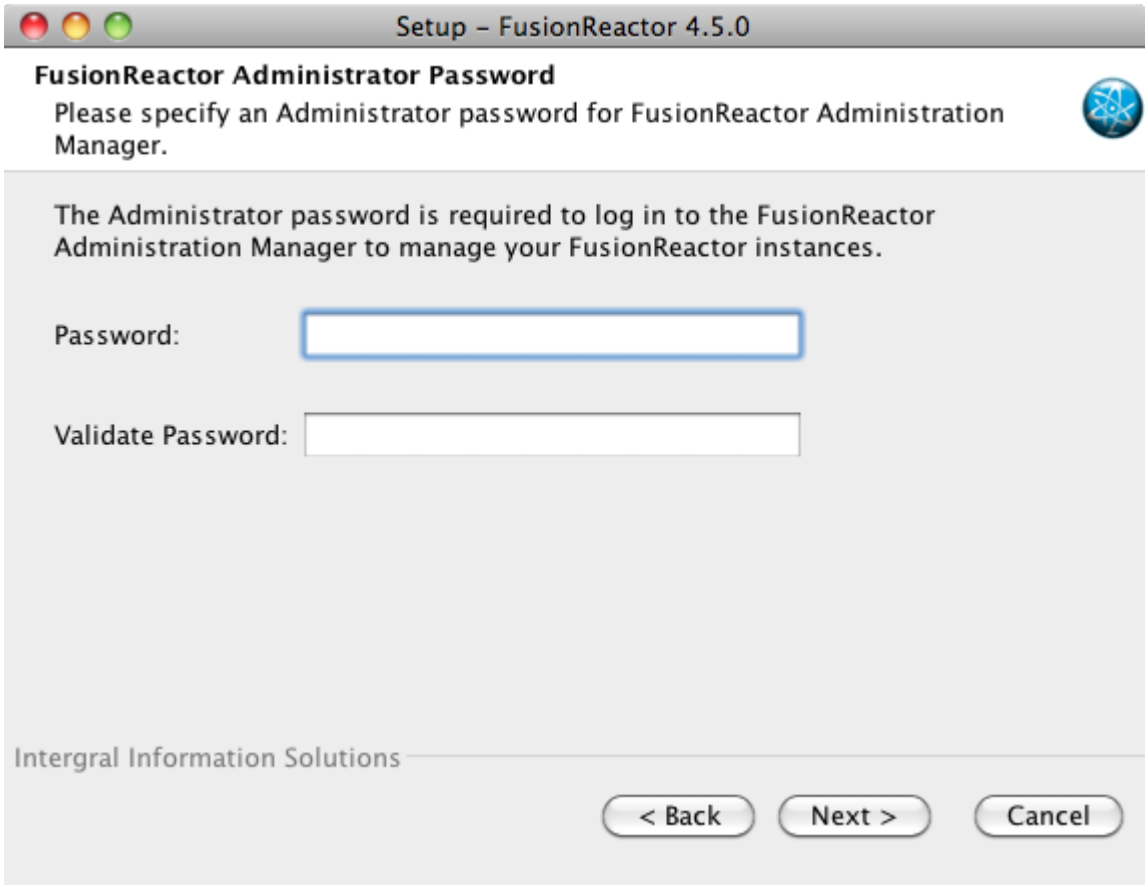


Figure 6: Administration Manager Password

Figure 6: Specify the password to be used by the Administrator user on FRAM. This password must be more than 5 characters long.

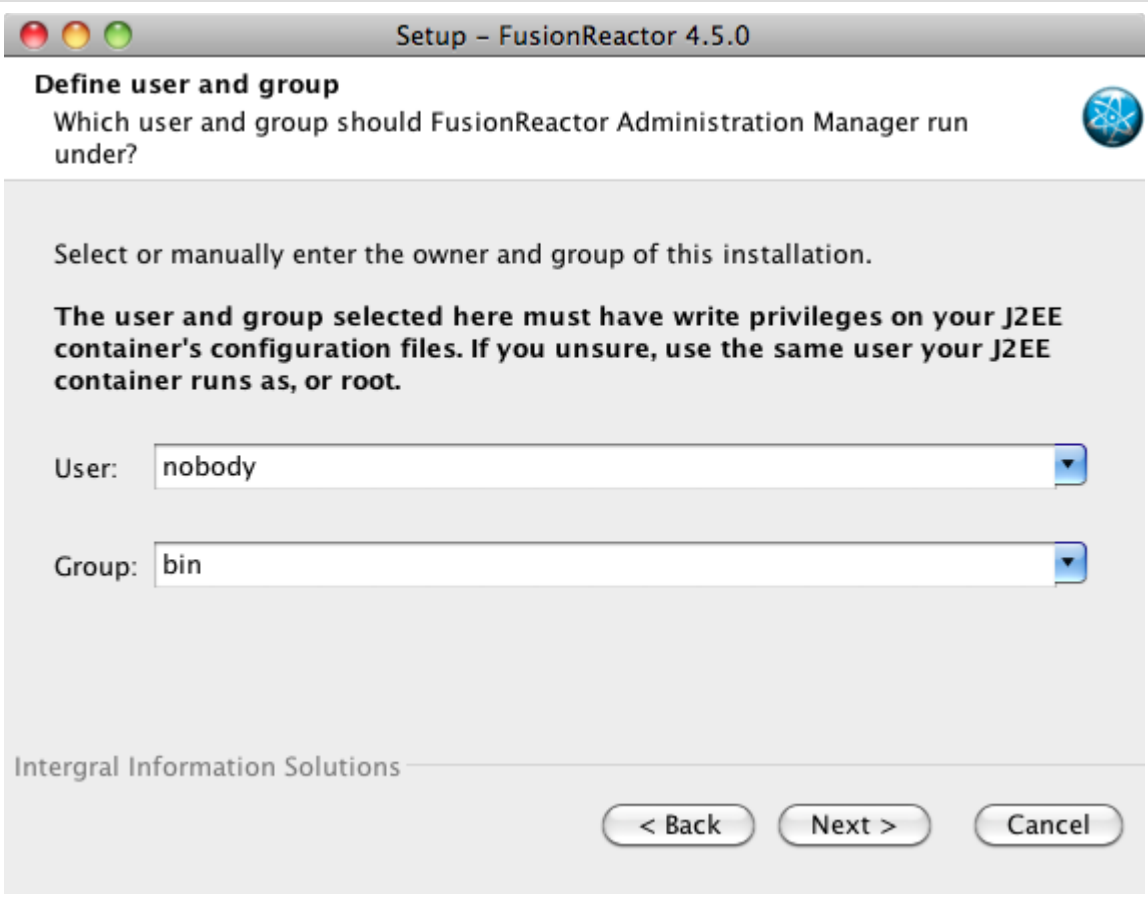


Figure 7: Choose a User and a Group

Figure 7: You will be asked to specify the owner and associated group of the FusionReactor installation directory (including FRAM). Only that user or someone with the same or higher privileges will be allowed to run FRAM.

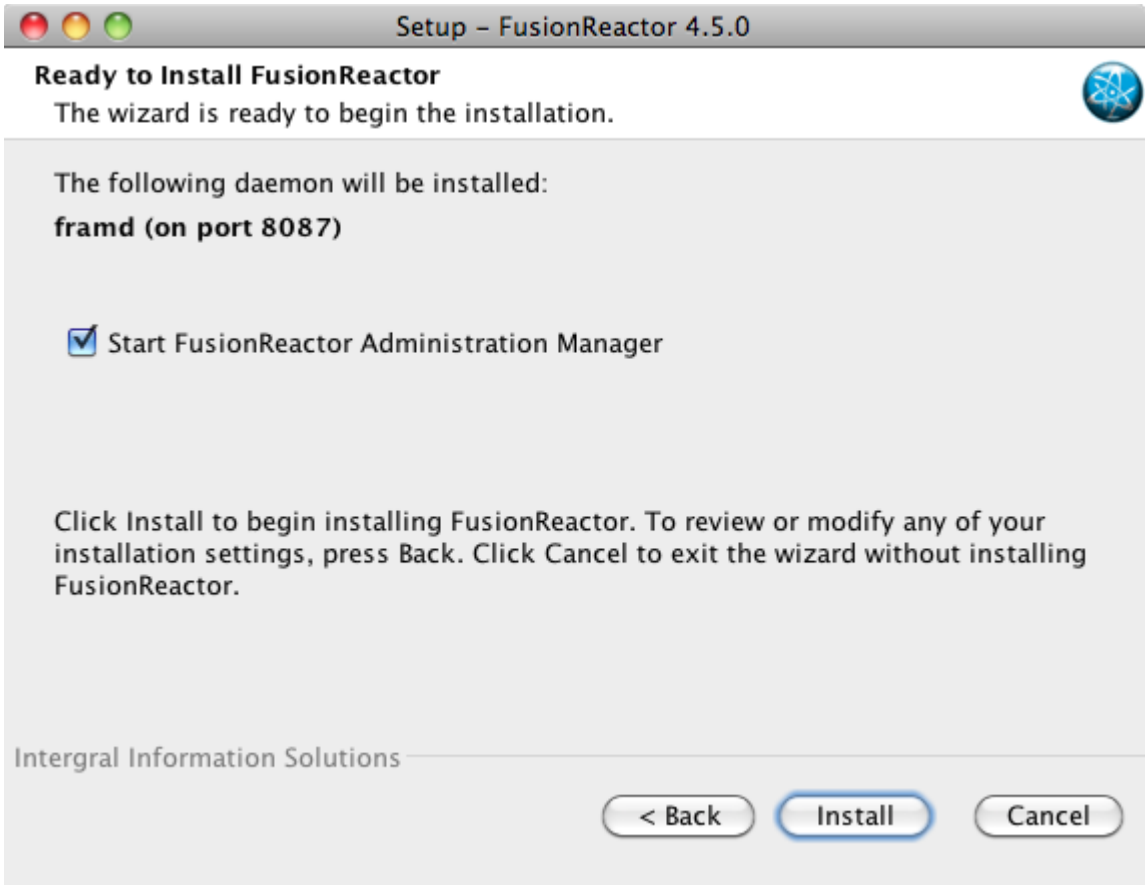


Figure 8: Ready to Install

Figure 8: Setup can automatically start FRAM as part of the installation procedure. If you do not wish to start FRAM at this moment, or wish to manually start FRAM uncheck the option *'Start FusionReactor Administration Manager".*

Click Install to begin updating to FusionReactor 4.5.x, Back to review the installation configuration or Cancel to exit the installation wizard.

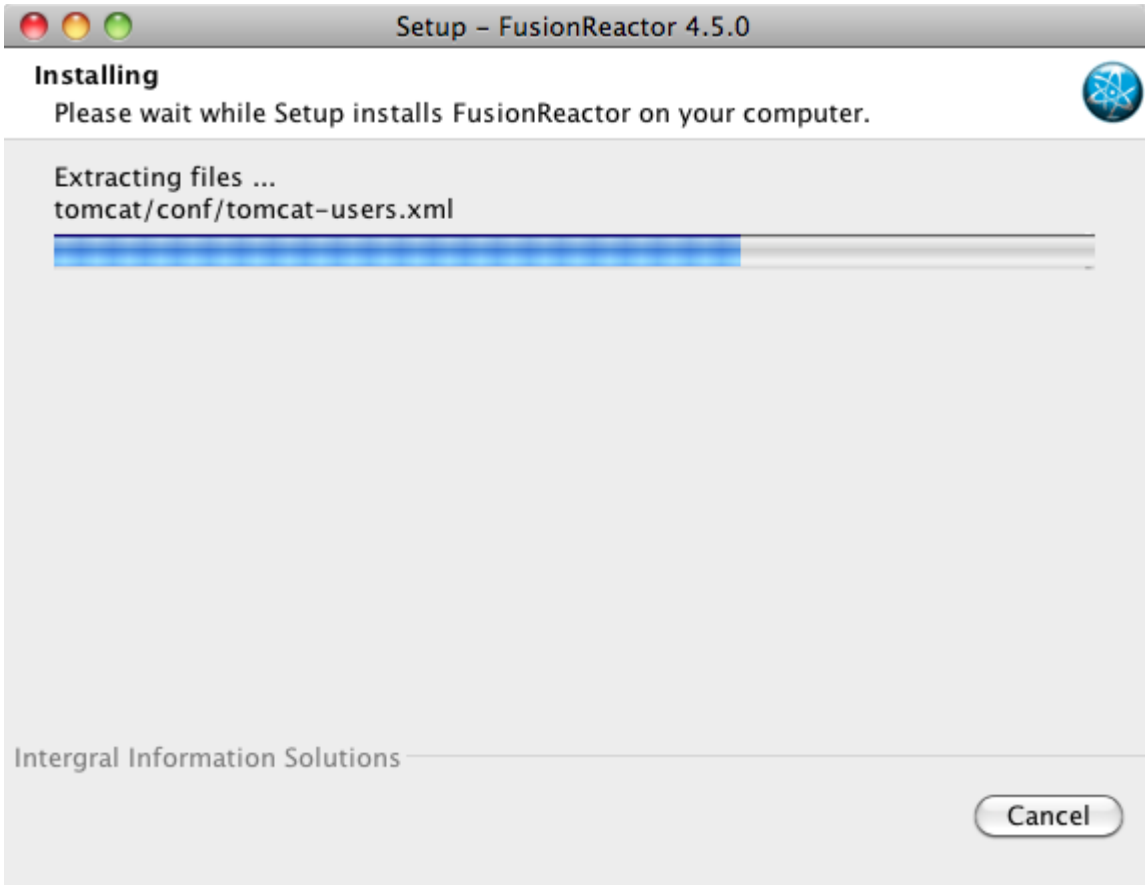


Figure 9: Installation Progress

Figure 9: The installation of FusionReactor 4.5.x may take a few minutes depending on the machine running the Setup.



Figure 10: Configure Instances Screen

Figure 10: After the installation has finished you are presented with a screen that allows you to decide if you want to configure the instances now or later. To configure now launch the Instance Manager by clicking the button 'Start configuring instances' and follow the Initial Server Scan Guide. If you wish to configure your instances later uncheck the option *'Open the Instance Manager on close"* and click the button 'Configure instances later'. You are only presented with this screen if you choose to start FusionReactor Administration Manager in Figure 8.

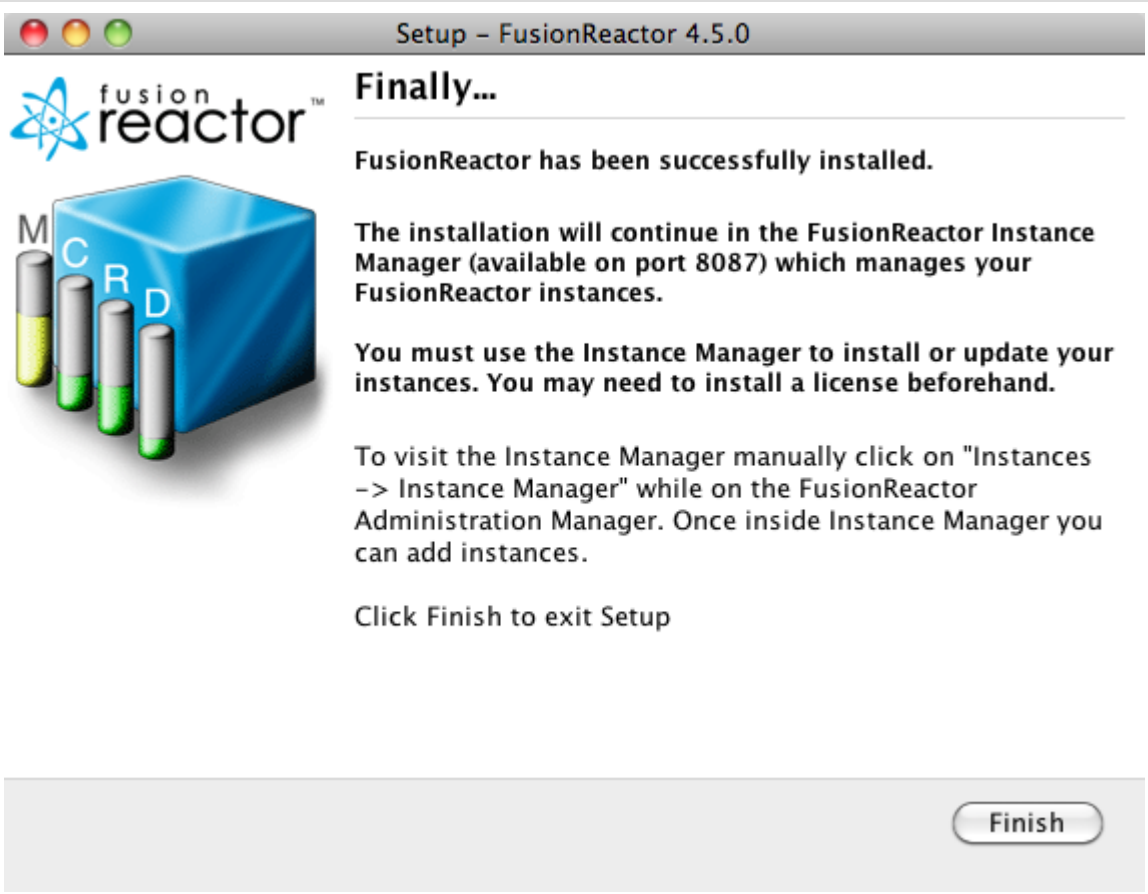


Figure 11: Final Screen (If you did not choose to start FusionReactor Administration Manager)

Figure 11: This screen provides you with a link to the FusionReactor Administration Manager and an overview on how to access Instance Manager as well as some additional information. You are only presented with this screen if you choose not to start FusionReactor Administration Manager in Figure 8. This is the last screen in the setup, clicking Finish will close the window.

Uninstalling FusionReactor on Mac

To completely remove FusionReactor from your system, you can either double click on the FusionReactor Uninstaller icon in the FusionReactor directory or you can run the uninstall command located in the FusionReactor directory in a shell window.

The Uninstaller will completely remove the FusionReactor directory including all of the generated data – **please take care that you save any files you want to keep before running the uninstaller.**

> The uninstaller will also remove any installed FusionReactor instances, during this process you will be asked to restart any running servers. Please note that not stopping the servers means that FusionReactor and all of its components will not be completely removed. This may cause issues on these servers.

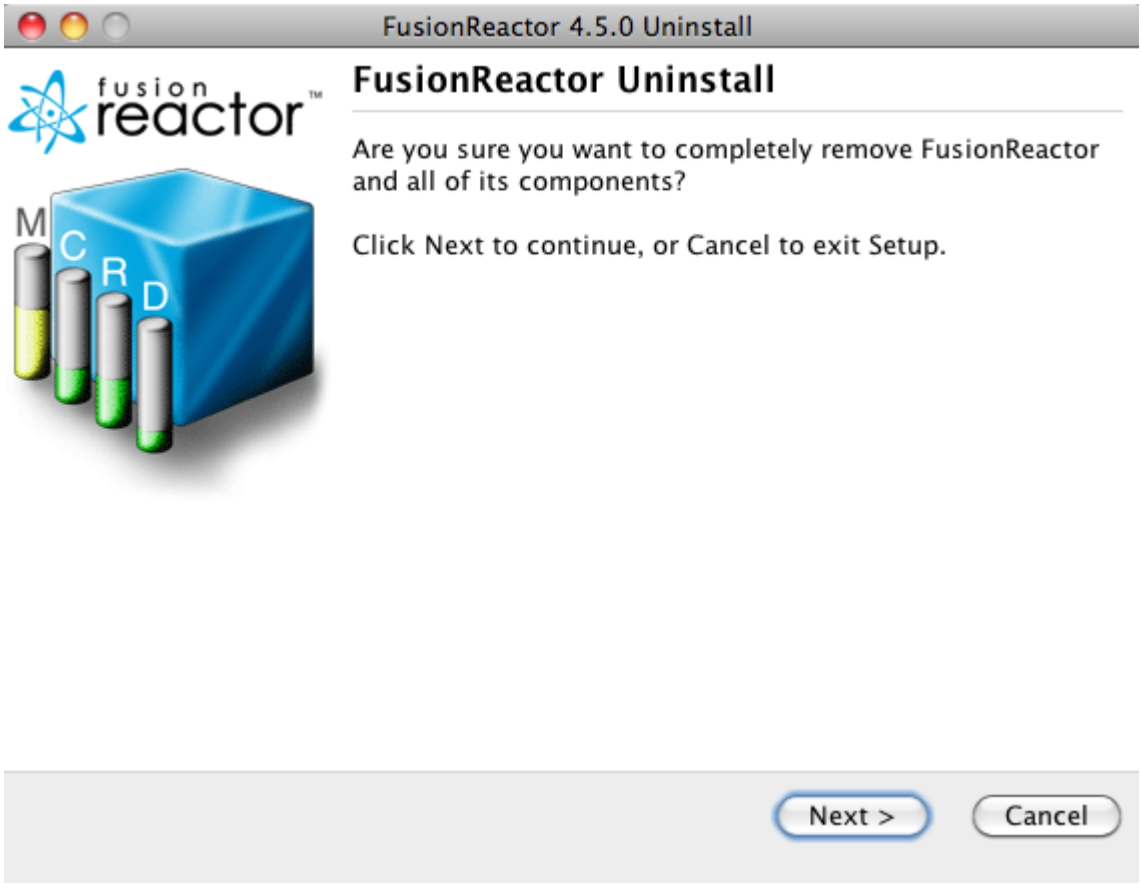


Figure 1: FusionReactor Uninstall

Figure 1: Once the Uninstaller has been started, read the text that asks you if you want to remove FusionReactor and all it's components. Click Next to continue on with the uninstall or click Cancel to exit the uninstaller wizard.

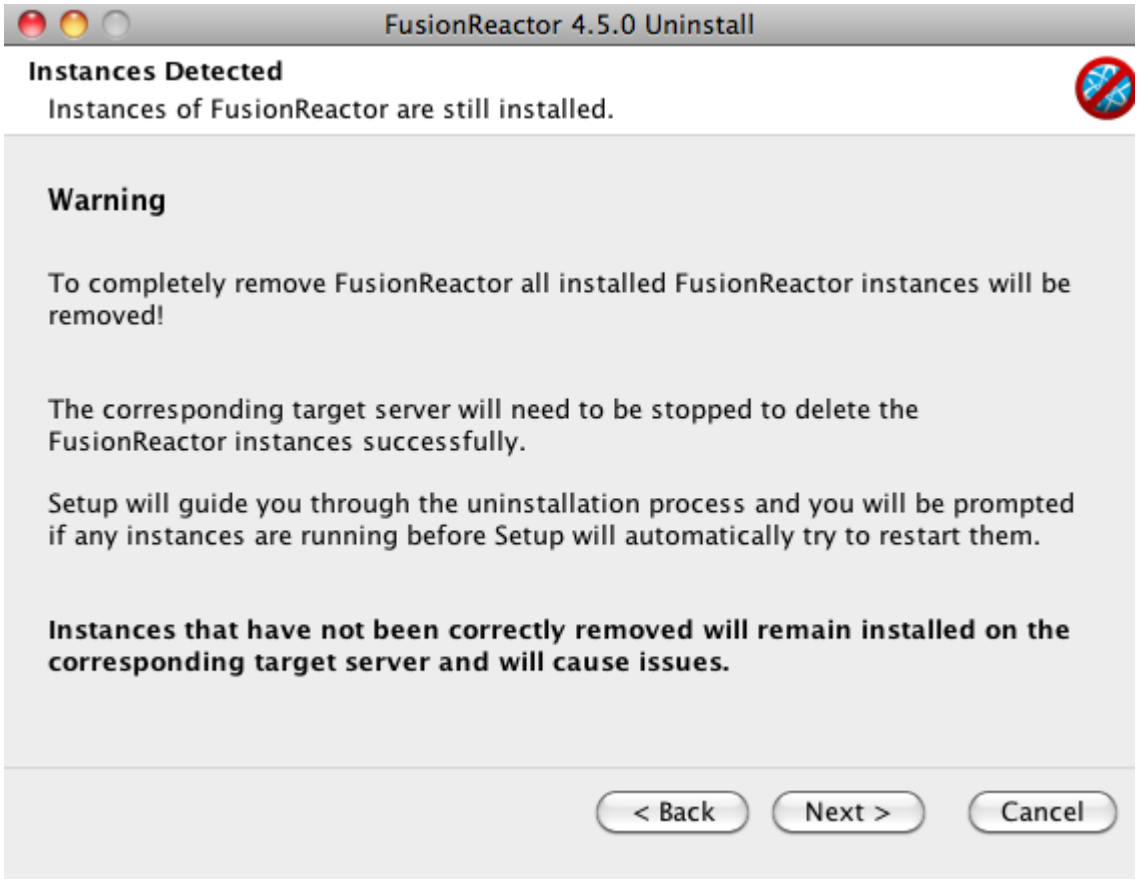


Figure 2: Uninstall FusionReactor Instances Warning

Figure 2: If you have other FusionReactor instances installed a warning message will be displayed to highlight the effects of removing FusionReactor. Carefully read through this. Once you have read it and you have understood it, please **click Next to uninstall all of the installed FusionReactor instances as well as FusionReactor.** Otherwise, click Cancel to stop the uninstall.

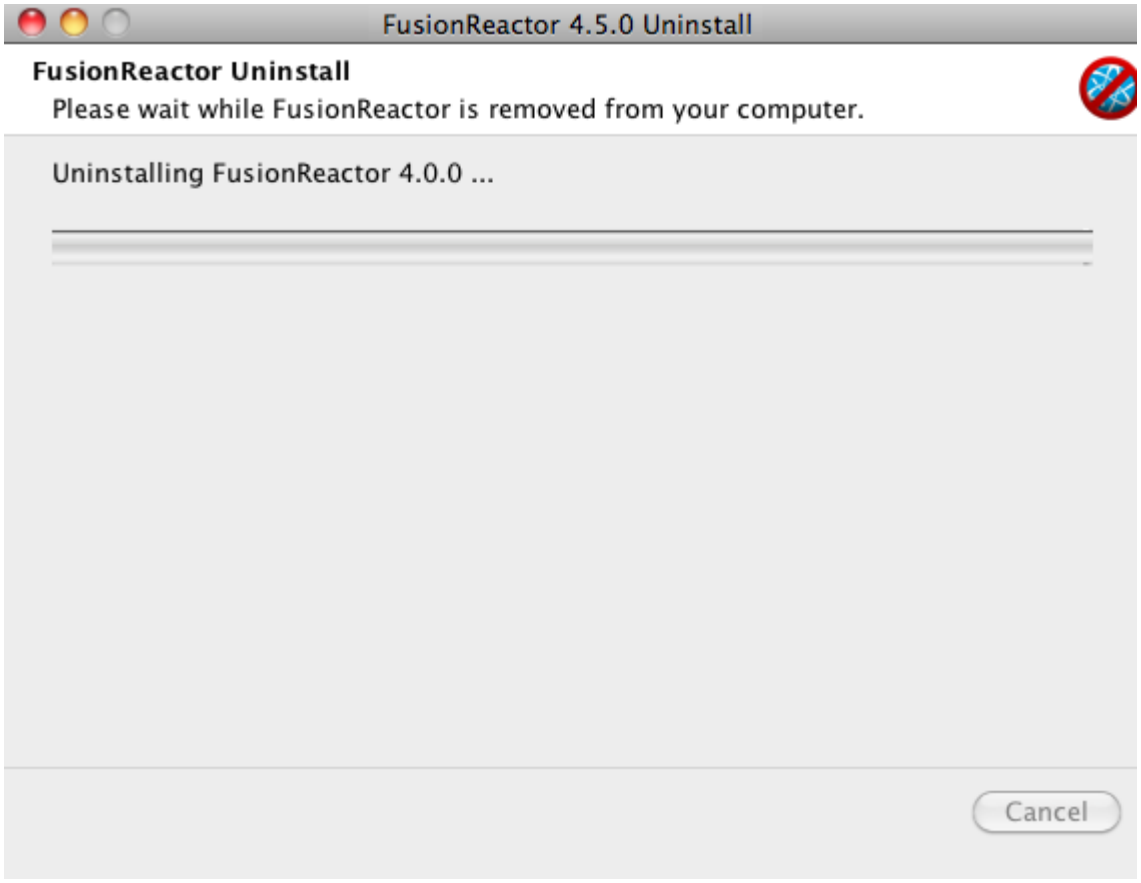


Figure 3: Uninstall Progress Bar

Figure 3: Once you have clicked Next, the FusionReactor uninstaller will begin to uninstall FusionReactor and all of the installed FusionReactor instances.

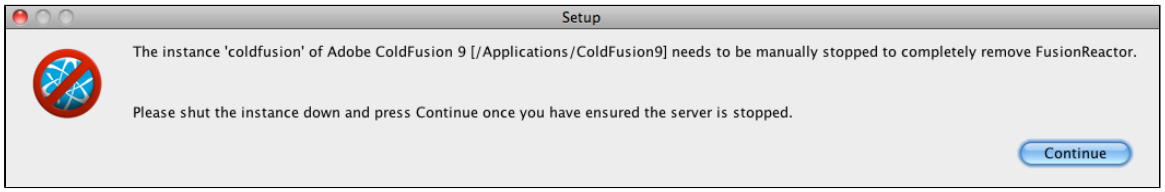


Figure 4: Stopping a Server Instance

Figure 4: If there is a FusionReactor instance that is running, the server container for this needs to be stopped in order to successfully remove the FusionReactor Instance.

Stop the server on which the FusionReactor instance is installed and press Continue once ensured the server is shut down, or press Continue if you do not wish to stop the server.

> As stated at the top of this page, not stopping the server container means that FusionReactor and all of its components will not be completely removed from the target server. This may cause issues on the corresponding server.

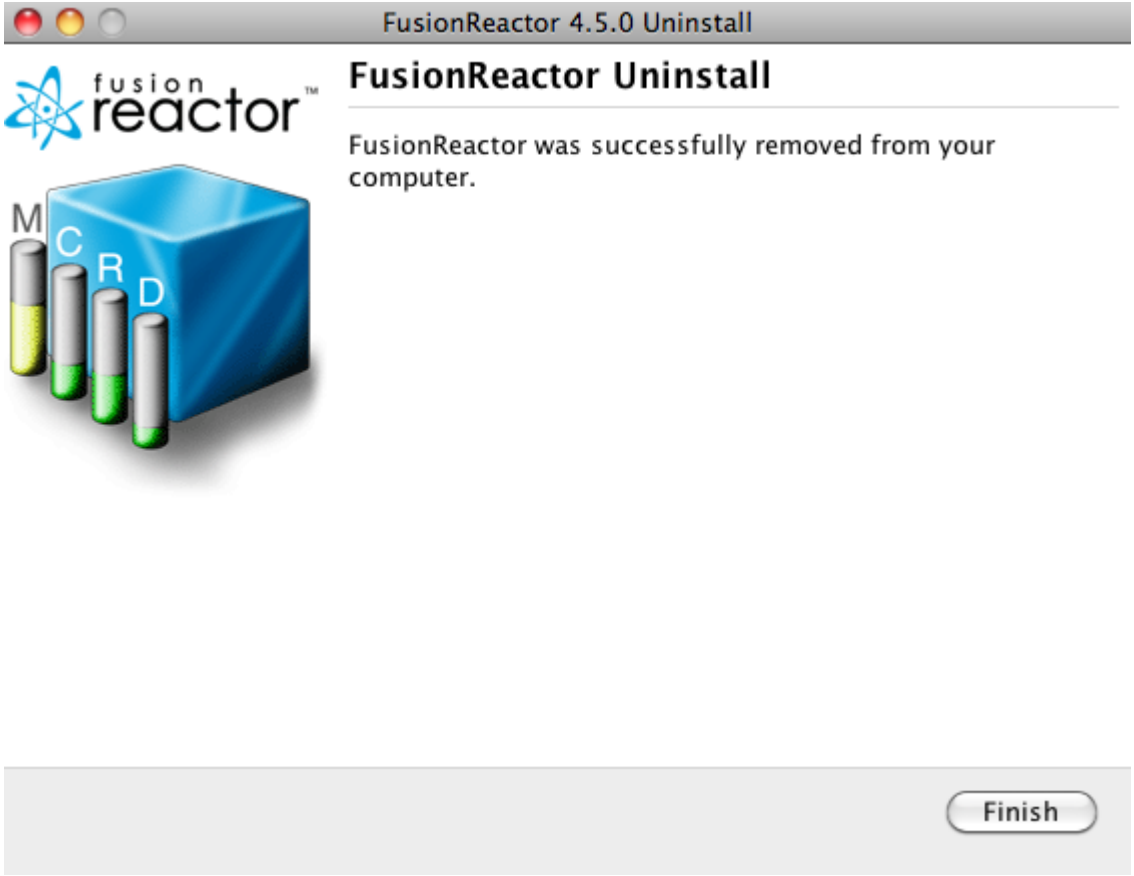


Figure 5: FusionReactor Uninstall Complete

Figure 5: When FusionReactor has finished uninstalling you will be presented with the Finish screen shown in Figure 6. Click on Finish to exit.

Uninstalling FusionReactor on Solaris

To completely remove FusionReactor from your system, you can either double click on the FusionReactor Uninstaller icon in the FusionReactor directory or you can run the uninstall command located in the FusionReactor directory in a shell window. Both methods require you to be logged in as the root user.

The Uninstaller will completely remove the FusionReactor directory including all of the generated data – **please take care that you save any files you want to keep before running the uninstaller.**

> The uninstaller will also remove any installed FusionReactor instances, during this process you will be asked to restart any running servers. Please note that not stopping the servers means that FusionReactor and all of its components will not be completely removed. This may cause issues on these servers.



Figure 1: FusionReactor Uninstall

Figure 1: Once the Uninstaller has been started, read the text that asks you if you want to remove FusionReactor and all it's components. Click Next to continue on with the uninstall or click Cancel to exit the uninstaller wizard.

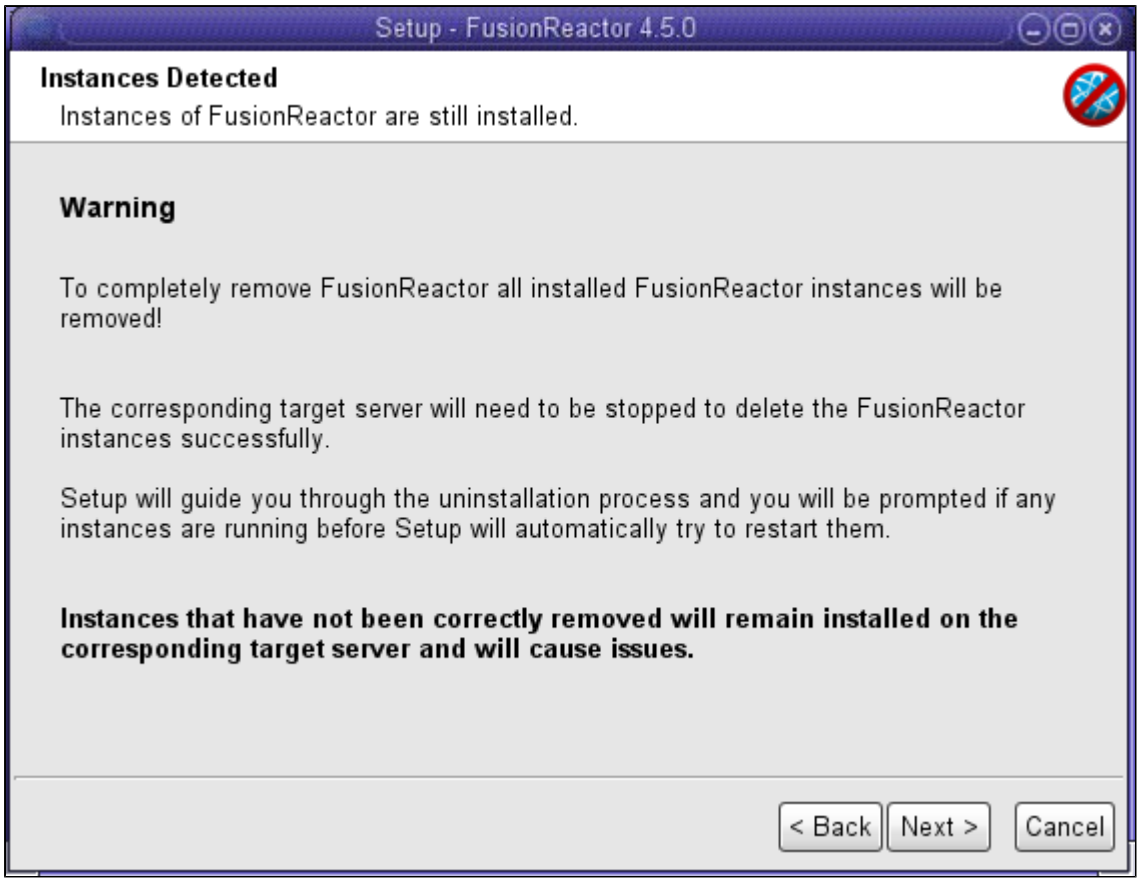


Figure 2: Uninstall FusionReactor Instances Warning

Figure 2: If you have other FusionReactor instances installed a warning message will be displayed to highlight the effects of removing FusionReactor. Carefully read through this. Once you have read it and you have understood it, please **click Next to uninstall all of the installed FusionReactor instances as well as FusionReactor.** Otherwise, click Cancel to stop the uninstall.

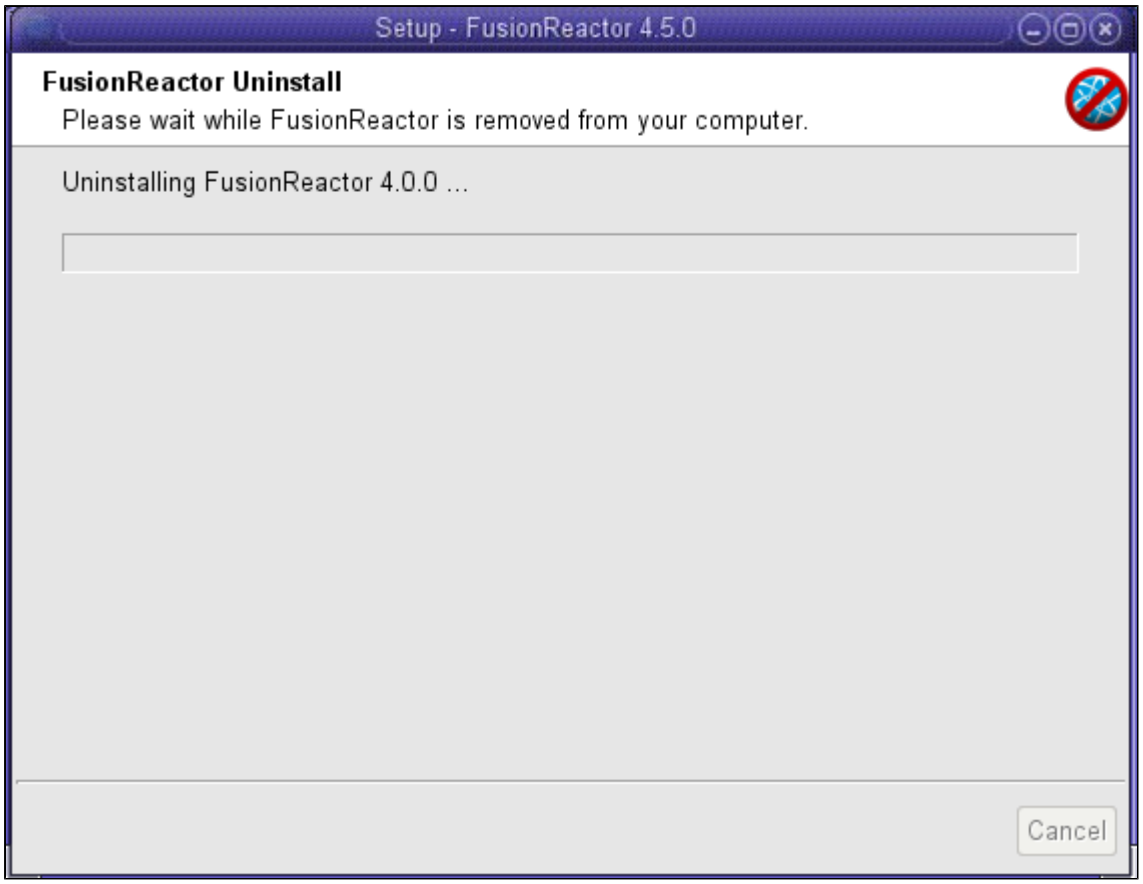


Figure 3: Uninstall Progress Bar

Figure 3: Once you have clicked Next, the FusionReactor uninstaller will begin to uninstall FusionReactor and all of the installed FusionReactor instances.

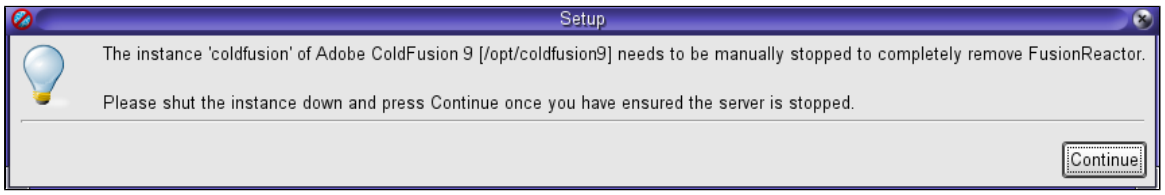


Figure 4: Stopping a Server Instance

Figure 4: If there is a FusionReactor instance that is running, the server container for this needs to be stopped in order to successfully remove the FusionReactor Instance.

Stop the server on which the FusionReactor instance is installed and press Continue once ensured the server is shut down, or press Continue if you do not wish to stop the server.

> As stated at the top of this page, not stopping the server container means that FusionReactor and all of its components will not be completely removed from the target server. This may cause issues on the corresponding server.

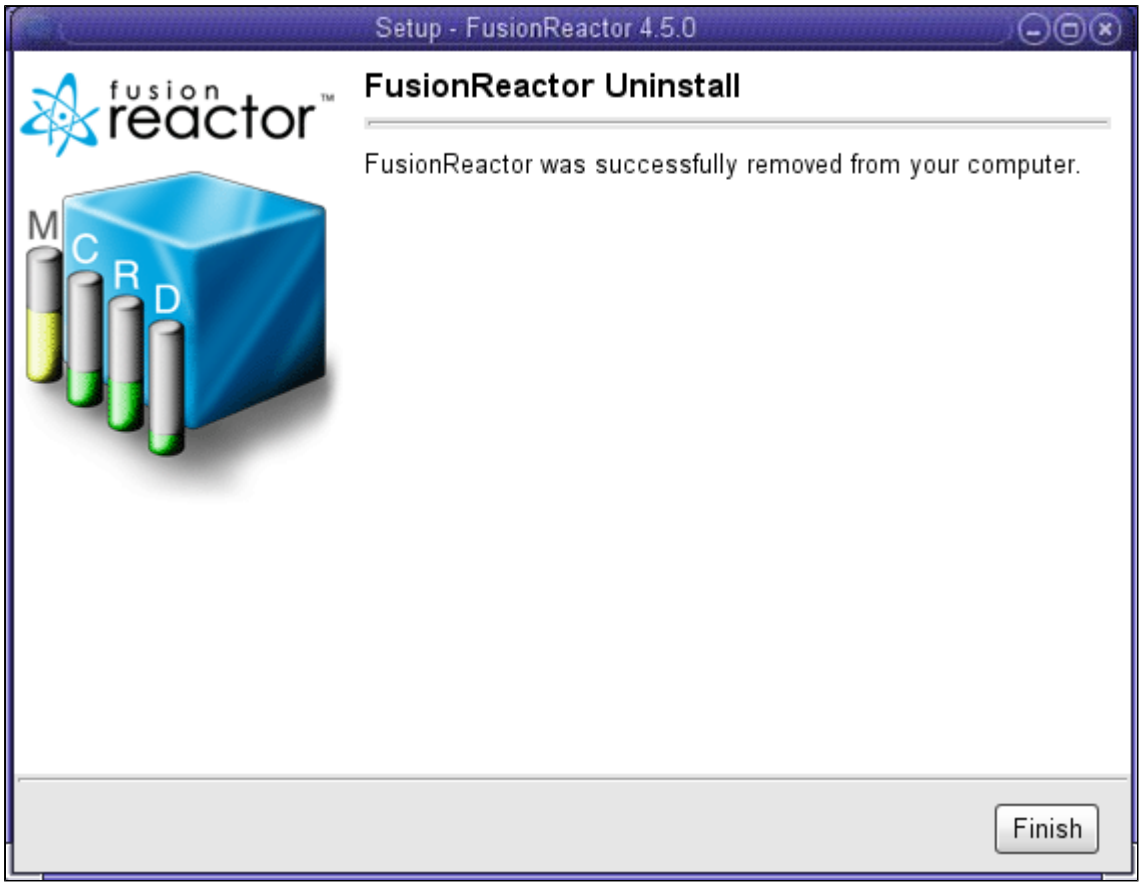


Figure 5: FusionReactor Uninstall Complete

Figure 5: When FusionReactor has finished uninstalling you will be presented with the Finish screen shown in Figure 6. Click on Finish to exit.

Where are FusionReactor 4's Manual Installer Packages?

> The manual install packages (Linux RPM, Windows ZIP, Generic Manual Installer and Generic Update) have been discontinued.

Because of the changes in structure with FusionReactor 4, including an extended folder structure for instances and more comprehensive configuration, it's no longer possible for us to supply manual installs or updates. The changes and customization required for each update would be too great to perform reliably by hand.

FusionReactor Administration Manager (FRAM)

Instead of installing FusionReactor directly into your J2EE container (e.g. ColdFusion), you now use FRAM to do this. FRAM, the lightweight Administration Manager, now takes care of individual installs in its **Instance Manager**, which is now available to both Enterprise and Standard customers.

No GUI

Most of our customers used the manual installer because they were running on Linux or Unix servers, and didn't have X-Windows available to provide a GUI.

Some customers used manual installs because our installer didn't recognize their J2EE server.

In the first case, we've added full command-line capability to the installer – just start it with the **`-c`** option in a terminal login - no GUI required.

```
sudo ./FusionReactor_linux_4_5_0.sh -c
```

In the second case, try the following:

- In FRAM's **Instance Manager**, have it scan exactly the folder in which your J2EE container is installed.
- If it's still not found, try to add it as a supported server using the **Add Server** form at the top of the Instance Manager.

If, after that, it's *still* not found...

Last-Resort Manual Install

Firstly, please get in touch with us. We'd love to know what server you're having trouble with so we can add it to a future release of FusionReactor. FusionReactor conforms to the J2EE Servlet/Filter specification, so it can be made to run on almost all J2EE containers - we just have to add the right code to the Instance Manager to detect them.

Secondly, there is a procedure for performing a manual install. If you are trying to update a pre-4.0 manually-installed version of FusionReactor, you'll have to uninstall it first, then get started with the procedure.

Manual Instance Installation

Manual Instance Installation

Introduction

The normal method for installing FusionReactor into a J2EE container is to use the **FusionReactor Administration Manager**'s **Instance Manager**. This page, located in **`Instances -> Instance Manager`**, allows you to easily administer all instances present on a single machine.

If FusionReactor doesn't detect your instance, you may have to install it manually. There are some caveats to this process:

- The instance will not be displayed within the **Instance Manager (IM)**.
  - You won't be able to use the IM to administer it. Internal webserver port, IP address and the "Bind to all interfaces" flag will only be changed by editing the `fram.conf` file (more on that later).
- The instance will be unable to use the FusionReactor Administration Manager to offload the zipping of logfiles, or the shipping of logfiles to **FusionAnalytics**.

Procedure

This procedure has been developed for Unix-type operating systems, including Linux. FusionReactor may be installed manually on windows, by changing the corresponding paths. A Windows-specific template configuration is available, and this is noted later in this procedure.

Prerequisites

The following list outlines the prerequisites which must be fulfilled prior to beginning this procedure.

- The target J2EE container must be installed and running.
- If you have previously installed FusionReactor 3.x manually, this must be uninstalled.
  - The FusionReactor filter definition, which is inserted usually in a `web.xml` (or `default-web.xml`) must be removed.
  - The FusionReactor installation folder (Unix: `/opt/fusionreactor`, Windows: `C:\FusionReac`

tor) must be removed.
- You should know in which file your J2EE server defines its servlets and filters. You'll need to know this in order to add the FusionReactor filter later.

This will give a clean environment in which FR4 can be installed.

Installing FusionReactor Administration Manager

The FusionReactor Administration Manager, or FRAM, is a cut-down version of FusionReactor which runs inside its own server, and provides administrative and monitoring functionality for other FusionReactor instances.

Although, as a manual install, you won't be able to take advantage of these functions, FRAM does contain an instance template which you can use to create a manually-installed instance.

- Unix users who do not have X-Windows or another GUI available to them can run the FusionReactor 4 installer in **command-line mode** (available in Windows, Linux and Solaris). Simply start it with something like `./FusionReactor_linux_4_5_0.sh -c`
- Unix users with a GUI available, and Windows users, can proceed through the graphical install wizard.

After the wizard completes, and assuming default answers, the FusionReactor Administration Manager will be installed into `/opt/fusionreactor` (`c:\FusionReactor` on Windows).

It is not strictly necessary to start FRAM, but if you wish to do so, it can be started using the shell script (`/opt/fusionreactor/tomcat/bin/framd start`), or by starting the **FusionReactor AM** service on Windows. If you picked the default port, the service will be available in your browser at `http://localhost:8087/`.

Installing a Manual Instance

FusionReactor Administration Manager includes a template instance which you can customize. The three main steps are:

- Creating a new instance folder and customizing the configuration.
- Inserting the filter definition (a small snippet of XML) into your J2EE container's `web.xml` (or `default-web.xml`) file.
- Copying the `fusionreactor.jar` file to a location accessible by your J2EE container instance.

Creating a New Instance Folder.

For the purposes of illustration, we will be installing a manual instance into Adobe's ColdFusion server, located in `/opt/coldfusion9`.

- Locate the supplied instance template folder. This will be in `/opt/fusionreactor/etc` and is called `instance-template`.
- Copy this template to `/opt/fusionreactor/instance`, and **rename it** to a unique name. This is the the **instance name**.
    - In our case, we're going to rename it to `cf9`.

Customizing the Instance

Within the `/opt/fusionreactor/instance/cf9/conf` folder, you will see two subfolders, `unix` and `windows`.

- Copy **either** the contents of the unix **or** the contents of the windows folder to `cf9/conf`.

You should now have a `reactor.conf` and a `filter-def.xml` file in the `conf` folder.

- Open the `reactor.conf` file with an editor.
- Search for all occurences of the string **__INSTANCENAME__** and replace them with the name you chose above.
    - In our case, this is cf9.

- If you have installed FusionReactor in a location other than the standard `/opt/fusionreactor` (`C:\FusionReactor`):
    - Search for all occurences of the installation folder (`/opt/fusionreactor`) and change them to your install location.

The internal web server is configured in the file `fram.conf`.

- Open the `fram.conf` file with an editor.
- The internal webserver, allowing access to FusionReactor over a port independent of your main webserver, is controlled using the **`webserver.*`** keys.

Key	Meaning
`webserver.active`	Whether the internal webserver is made available. If not, FusionReactor will be available over your main webserver under `http://localhost/fusionreactor`.
`webserver.address`	The IP address to which the internal webserver should bind.
`webserver.all`	Whether the webserver should bind to **all** addresses.
`webserver.port`	The port on which the webserver listens.

Perhaps most important is the last key: `webserver.port`. You must ensure you use a port here which is not in use by any other application.

Adding the Filter Definition

The file `filter-def.xml` contains an XML snippet which must be added to your J2EE container's `web.xml` file. You should ascertain the location of this file by referring to the container's documentation.

- Open the `filter-def.xml` file.
- Copy the whole content of this file (a `filter` block and a `filter-mapping` block) into the clipboard.
- Open the filter/servlet definition file.
    - In ColdFusion's case, this is `/opt/coldfusion9/runtime/servers/coldfusion/SERVER-INF/default-web.xml`.
- Below the `web-app` start tag, just before the first `filter` definition (if one is present, if not, just before the first `servlet` definition), **paste** the complete content of the clipboard.
- Customize the `__INSTANCENAME__` placeholder to point to your newly-created instance folder.

You should now have inserted the FusionReactor `filter` and `filter-mapping` block just prior to the first filter (or servlet) definition.

Copying the `fusionreactor.jar` file

Within the folder `/opt/fusionreactor/etc/lib`, you'll find a copy of the `fusionreactor.jar` file. This must be copied to somewhere visible to your J2EE server's webapp instance.

* In our case, we copy the file to `/opt/coldfusion9/runtime/servers/coldfusion/SERVER-INF/lib/`.

Finally, start your J2EE container, and watch its standard-out and standard-error logs. If anything is

misconfigured, you should be able to start debugging the problem from the output presented there.

Installing FusionReactor in Locked Down Environments

Introduction

This page shows you how to get FusionReactor up and running in environments where your J2EE server is running as a different user than FusionReactor's FRAM service. The procedure is different depending on your operating system, so we've split it into sections.

What does "locked down" mean?

In order to help secure their production operating systems, many people choose to run their J2EE server *not* as root (or Local System on windows), but as a normal user, "locking down" the application.

This is done to limit the effects of an attacker using malicious exploit to gain access to the server. An attacker who exploits a weakness in a running server usually attempts to have that server run his arbitrary code, essentially gaining access to the system. By running the server under a normal user account, rather than root (or Local System), the attacker can only gain the privileges associated with the normal user, rather than those of the superuser.

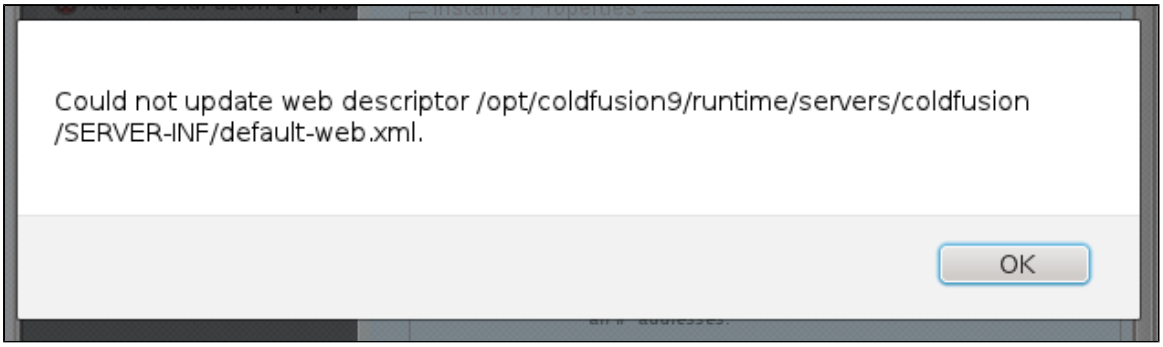
It is considered good practice to separate the privileges of servers and daemons by running them as normal users in this way, thereby granting only the minimum privilege needed to operate correctly.

How does this affect FusionReactor?

While we encourage users to lock down their J2EE systems, this causes two issues for FusionReactor.

1. **FusionReactor's Administration Manager** (FRAM, usually on port 8087) requires write access to parts of the J2EE server in order to add itself as a monitoring element.
2. Each **FusionReactor** instance in an installation requires read/write access to the **FusionReactor** installat ion folder, in order to - for example - serve HTML objects, write persistent configuration and so on.

In a secure environment, FusionReactor may be running as one user, while the J2EE server runs as another. This means that users may - for instance - see the following message when trying to install an instance:



FusionReactor was not able to modify the J2EE web descriptor file because it's owned by a different user than that running the FusionReactor Administration Manager.

Solution

The solution is to create a common group, add both users to this group, and use File Access Control Lists (FACLs) to grant access to both FusionReactor and the J2EE server to this group. There is relatively low risk in this solution, as it does not involve granting "world" (Everybody) permissions. Only the two users will be able to access each others' resources.

We have developed the following procedures for Linux, Solaris and Windows. The procedure does the

following:

1. Create a new group "fusion" to be a common group of both the FusionReactor runtime user ("nobody" on Linux) and the J2EE server user ("coldfusion").
2. Add both users to this group.
3. Sets file ACLs and Default ACLs (so new files and directories inherit the ACL rights) on the J2EE server and FusionReactor.

The procedure uses Adobe's ColdFusion (JRun) J2EE server for its examples, with FusionReactor running as "nobody" on Linux/Solaris. If you have multiple J2EE servers, their users must also be added to the common group, and their Access Control Entries (ACE) and default ACEs must also be set.

If you later add new users and J2EE servers, simply add their users to the common group, and add Access Control Entries as detailed in each procedure.

Linux

Requires kernel 2.6 or later, with a filesystem supporting ACLs (ReiserFS, ext2, ext3, JFS, XFS), mounted with ACL support (default).

On Linux, FusionReactor is typically installed as **nobody**. In our test environment, ColdFusion is installed as the **coldfusion** user. If your users differ, substitute them in the commands below. The following commands should be run as **root** with FusionReactor (framd) **and** your J2EE server stopped.

> ⚠ If your **nobody** user has a shell which prevents login (e.g. **/sbin/nologin**), you may need **firs t** apply technote FRS-284 to resolve this issue.

1. Add the common group **fusion**

```
groupadd fusion
```

2. Add FusionReactor's runtime user, and ColdFusion's (substitute your J2EE server's) runtime user to the common group

```
gpasswd fusion --members nobody,coldfusion
```

3. Add file Access Control Entries to both FusionReactor and ColdFusion (substitute your J2EE server) directories

```
setfacl -R -m g:fusion:rwx /opt/coldfusion9 /opt/fusionreactor
```

4. Add default file Access Control Entries to the same directories. This allows either system to create new directories and have the entries propagate.

```
setfacl -R -m d:g:fusion:rwx /opt/coldfusion9 /opt/fusionreactor
```

Windows

Requires Windows Server 2003 Service Pack 2, Windows Vista, Windows Server 2008 or Windows 7.

On Windows, FusionReactor is installed as the **Local System** user. Running as another user is currently

unsupported. In our test environment, we are running ColdFusion (our J2EE server) as user **coldfusion**. The following commands should be run in a **cmd** window as a user who has **Administrator** rights to the system. If you cannot log in as an Administrator, a command window may be opened from a normal **cmd** session using the following command:

```
runas /user:Administrator cmd
```

FusionReactor **and** your J2EE server should be stopped.

1. Add the group **fusion**

   ```
   net localgroup /add fusion
   ```

2. Add ColdFusion's (substitute your J2EE server's) runtime user to the group

   ```
   net localgroup fusion /add coldfusion
   ```

3. Add a default file Access Control Entry to the ColdFusion (substitute your J2EE server) directory (this will propagate down automatically)

   ```
   icacls c:\ColdFusion9 /grant fusion:(OI)(CI)F
   ```

4. Add a default file Access Control Entry to the FusionReactor directory (this will propagate down automatically)

   ```
   icacls c:\FusionReactor /grant fusion:(OI)(CI)F
   ```

Solaris

Tested on Solaris SunOS 5.10 (Solaris 10).

On Solaris, FusionReactor is typically installed as **nobody**. In our test environment, ColdFusion is installed as the **cfusion *user. If your users differ, substitute them in the commands below. The following commands should be run as *root** with FusionReactor (framd) **and** your J2EE server stopped.

> ⚠ If your **nobody** user has a shell which prevents login (e.g. **/sbin/nologin**), you may need **firs t** apply technote FRS-284 to resolve this issue.

1. Add the group **fusion**

   ```
   groupadd fusion
   ```

2. Add FusionReactor's runtime user, and ColdFusion's (substitute your J2EE server's) runtime user to the common group

```
usermod -G fusion nobody
usermod -G fusion cfusion
```

3. Add file Access Control Entries to files in ColdFusion (substitute your J2EE server's directory) and FusionReactor. This must be done with **find** as Solaris 10 has no recursive option to **setfacl**.

```
find /opt/fusionreactor -exec setfacl -m g:fusion:rwx,m:rwx {} \;
find /opt/coldfusion9 -exec setfacl -m g:fusion:rwx,m:rwx {} \;
```

4. Add default file to the directories of ColdFusion (substitute your J2EE server) and FusionReactor. Again, this is done with **find**.

```
find /opt/fusionreactor -type d -exec setfacl -m
d:u::rwx,d:g::r-x,d:o:r-x,d:m:rwx,d:g:fusion:rwx {} \;
find /opt/coldfusion9 -type d -exec setfacl -m
d:u::rwx,d:g::r-x,d:o:r-x,d:m:rwx,d:g:fusion:rwx {} \;
```

JDBC Driver Wrapper User Guide

JDBC Driver Wrapper Overview

Installation in ColdFusion

Creating a New Wrapped Datasource in ColdFusion

- Using Macromedia's (DataDirect) ColdFusion Built-in Drivers
- Constructing JDBC URLs for other DataDirect Drivers (CFMX,7,8,9,10)
- Using a User-Specified Driver (“other”) Driver

Using the FusionReactor JDBC Wrapper

Interpreting JDBC Log Data

Prepared Statement Positional Bind Parameters Replacement Strings

A Note On SQL Server Select Methods

Exception Catalog

FusionReactor 4.5.0: Upgrading to 4.5.0 from the Split-Jar Procedure

JDBC Driver Wrapper Overview

The FusionReactor JDBC Driver Wrapper allows developers and administrators to control the interaction between Java and a database. The driver wrapper allows fine-grained metrics and reporting of database activity:

- *Logging of statements which ran against a database*
  This feature is useful to help detect deadlocks, see exactly what Statements look like without resorting to manual log output, and to see exactly how your Prepared Statements were bound before being run

against the database.

- *Row Limiting*
  The integrated row limiter can automatically halt database read activity after a user-specifiable number of rows is reached. This can stop run-away queries before they become a memory and resource problem.
- *Notification and Reminders*
  The driver can optionally notify you when a certain number of rows has been read, and periodically thereafter. Using this feature, you are able to keep a clear overview about the volume of data being processed by Java.

The driver wraps any existing JDBC driver and is able to communicate metric data to FusionReactor for easy perusal in the FusionReactor Administrator. If the driver does not detect FusionReactor running (for instance when running in a standalone Java application), it reports metrics to the standard output stream.

Intended Audience

This technical documentation is targeted at developers and administrators of standalone and J2EE/ColdFusion applications. It presents the procedure for installing and configuring the FusionReactor JDBC Driver Wrapper to run as a ColdFusion Data Source or as a standalone JDBC data source. The administrator is expected to have experience with ColdFusion Data Sources, and with how JDBC URLs are constructed.

Limitations

The FusionReactor JDBC Driver Wrapper is intended to help you debug and manage J2EE JDBC queries. It integrates tightly with FusionReactor to provide you with metrics and detailed information about your pages' database activities. Our JDBC Driver Wrapper can't optimize your queries and pages before they're run - but it *c an* help you see where the time is being spent, or to locate a 'stuck' query. The FusionReactor JDBC Driver Wrapper also can't totally insulate you from performance or stability issues with underlying drivers.

Some of our customers have experienced stability issues with the standard drivers shipped with ColdFusion, for example, and while FusionReactor helped to pinpoint the problem, it was ultimately resolved by trying an alternative driver. For example, many database vendors ship their own JDBC drivers, and they are often very satisfactory for production use.

Installation in ColdFusion

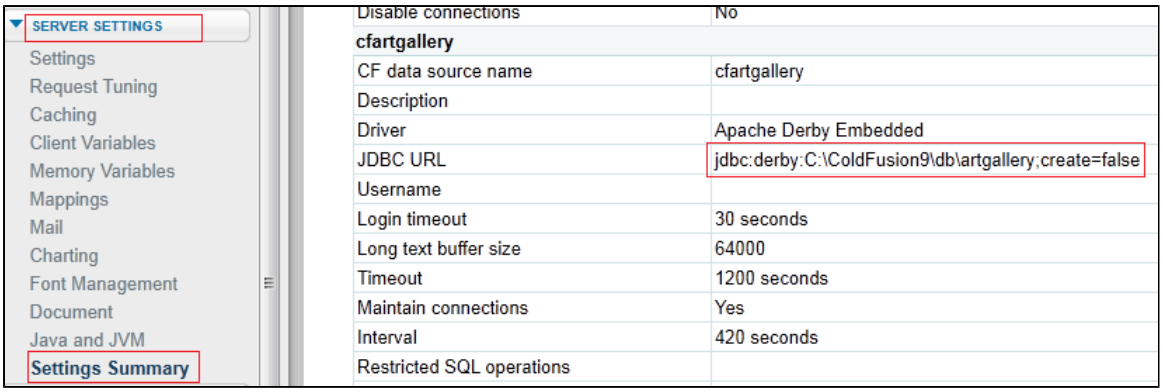
The FusionReactor JDBC Driver Wrapper is delivered as an integral part of the FusionReactor product, and is located within the fusionreactor.jar file, which is installed during the FusionReactor setup process. In order to use the driver with ColdFusion applications, you need only to alter the Data Source definition within the ColdFusion administrator.

Wrapping or Adding a New Datasource?

If you are wrapping an existing datasource, you should continue to follow the instructions in this section. If you are starting from scratch, skip ahead to “Creating a New Wrapped Datasource in ColdFusion”. This guide walks you through locating the JDBC URL of your data source in ColdFusion and then wrapping them using the ColdFusion wrapper.

Locating your Macromedia's (DataDirect) ColdFusion Built-in Drivers

In the ColdFusion Administrator click on **Server Settings** in the table of contents, then click on the **Settings Summary** link. Locate the section with the heading Database & Services and the table with the name Database Data Sources. Locate the JDBC URL of the datasource that you want to wrap with a FusionReactor wrapper and then copy it or write it down.



Your JDBC URL will most likely be different, the purpose of the screen shot is to give you an idea of what to expect in the ColdFusion User Interface. For the purpose of this example our JDBC URL will read:

```
jdbc:macromedia:sqlserver://int0006:1433;databaseName=frtest;sendStringParametersAsUnicode=false;MaxPooledStatements=1000
```

Wrapping your JDBC URL with FusionReactor

1. Click **Data and Services** and then click on the **Data Sources** link.
2. Add a data source name.
   This page walks you through locating your Macromedia (DataDirect) ColdFusion Drivers and wrapping them in the FusionReactor wrapper. The page is broken down into 3 sections, the first walks you through finding your JDBC URL for the data source, the second walks you through wrapping it and the last one gives you an example.
3. Select 'other'-type driver.
4. Click on the Add button, you will be presented with more text input fields.
5. In the field titled CF Data Source Name enter a name for your data source.
6. In the JDBC URL text field simply enclose the existing JDBC URL within the FusionReactor syntax:

   ```
   jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://int0006:1433;databaseName=frtest;sendStringParametersAsUnicode=false;MaxPooledStatements=1000}
   ```

   You can also give the data source a name. If you name the data source in this way, FusionReactor will display this name in the Data Source column of the JDBC tab in the Request Details page. This is useful if you're using multiple data sources. Below is an example:

   ```
   jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://int0006:1433;databaseName=frtest;sendStringParametersAsUnicode=false;MaxPooledStatements=1000};name=TestName
   ```

7. In the JDBC Driver Class field, enter the name of the FusionReactor JDBC Driver Wrapper class:

   ```
   com.intergral.fusionreactor.jdbc.Wrapper
   ```

8. Enter 'FusionReactor' in the Driver Name field.
9. Select an appropriate username and password values in their respective fields.
10. Click 'Submit'. ColdFusion will immediately test the driver and provide feedback if necessary. When the FusionReactor driver is loaded, it will report “FusionReactor JDBC: Driver loaded.” to the standard output stream (usually logged to coldfusion-out.log or default-out.log within cfmx\runtime\logs)

Below are two screen shots, figure 1 covers stages 1 to 4 and figure 2 covers stages 5 to 11.

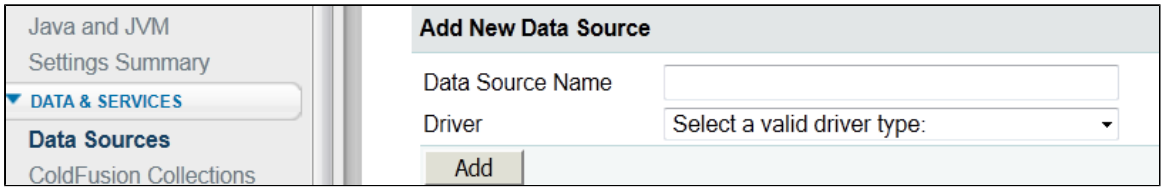


Figure 1

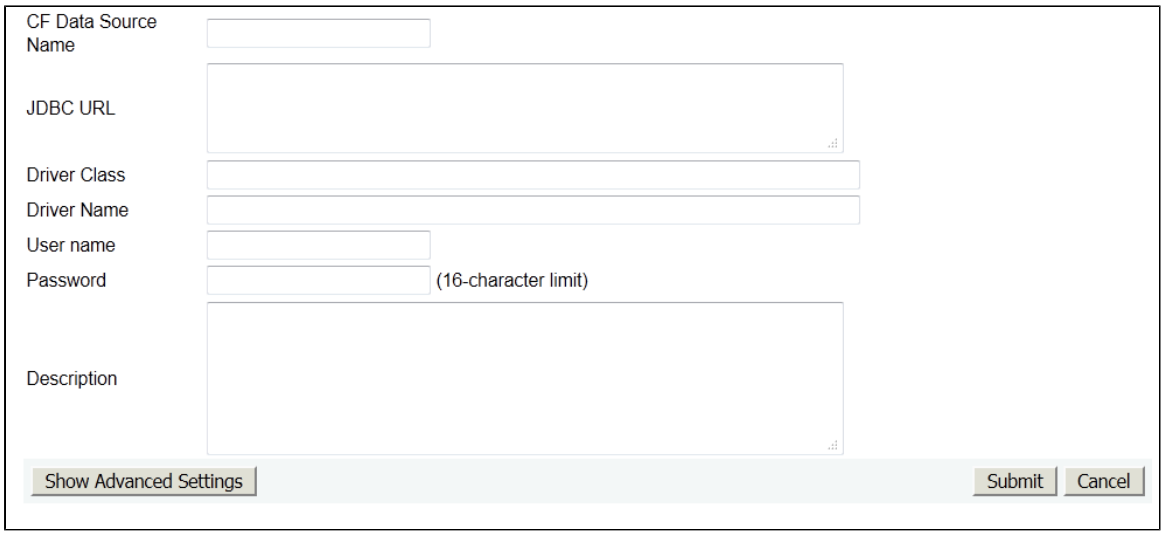


Figure 2

After submitting this form, ColdFusion will immediately test the connection and provide any necessary feedback in the ColdFusion Administrator. If the FusionReactor JDBC Driver Wrapper detects any problems with the underlying wrapped driver or with its own options, it will also provide feedback in the same way. Any exceptions generated by the FusionReactor JDBC Driver Wrapper are prefixed with an exception ID number, which may be used to look up more information. When running as part of J2EE/ColdFusion, and if the FusionReactor JDBC Driver Wrapper detects a running FusionReactor application, any queries run during the course of the request will be reported to FusionReactor for inclusion in the user interface, where appropriate.

Examples

The original data source was an MS-SQL Server with the database database name'testdb', after being wrapped by the FusionReactor wrapped it is now called 'MyDataSource'.

Original JDBC URL:

```
jdbc:macromedia:sqlserver://localhost:1433
;databaseName=testdb
```

FusionReactor URL:

```
jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://localhost:1433;databaseName=testdb};name=MyDataSource
```

Also See

Creating a New Wrapped Datasource in ColdFusion

Creating a New Wrapped Datasource in ColdFusion

This section shows you how to create a wrapped datasource from scratch. It explains how to construct JDBC URLs for ColdFusion's builtin drivers (called the Macromedia drivers, since they were introduced by Macromedia in ColdFusion MX), as well as for existing custom (non-Macromedia) datasources.

About Wrapped Datasources

Underlying all ColdFusion (and indeed all Java) database interaction is a configuration string called a **JDBC URL**. This string tells ColdFusion almost everything it needs to connect to the database: which driver to use, which machine to connect to, which port, and sometimes also the database login and password too. When using a standard ColdFusion driver, the ColdFusion Administrator application constructs this URL 'behind the scenes', using the details you enter into the form about your server. The FusionReactor JDBC Driver Wrapper requires this URL in order to function. Fortunately, finding it is quite simple.

Using Macromedia's (DataDirect) ColdFusion Built-in Drivers

Constructing JDBC URLs for other DataDirect Drivers (CFMX,7,8,9,10)

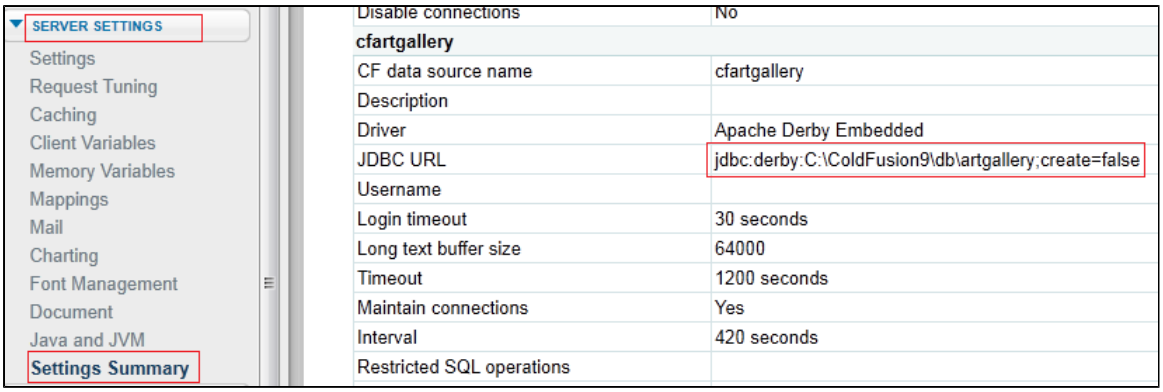
Using a User-Specified Driver (“other”) Driver

Using Macromedia's (DataDirect) ColdFusion Built-in Drivers

This page walks you through locating your Macromedia (DataDirect) ColdFusion Drivers and wrapping them in the FusionReactor wrapper. The page is broken down into 3 sections, the first walks you through finding the JDBC URL for your target data source, the second walks you through wrapping it and the last one gives you an example.

Locating your Macromedia's (DataDirect) ColdFusion Built-in Drivers

In the ColdFusion Administrator click on **Server Settings** in the table of contents, then click on the **Settings Summary** link. Locate the section with the heading Database & Services and the table with the name Database Data Sources. Locate the JDBC URL of the datasource that you want to wrap with a FusionReactor wrapper and then copy it or write it down.



Your JDBC URL will most likely be different, the purpose of the screen shot is to give you an idea of what to expect in the ColdFusion User Interface.

Wrapping your JDBC URL with FusionReactor

1. Click **Data and Services** and then click on the **Data Sources** link.
2. Add a data source name.
3. Select 'other'-type driver.
4. Click on the Add button, you will be presented with more text input fields.
5. In the field titled CF Data Source Name enter a name for your data source.
6. In the JDBC URL text field simply enclose the existing JDBC URL within the FusionReactor syntax:

   ```
   jdbc:fusionreactor:wrapper:{JDBC URL}
   ```

   You can also give the data source a name. If you name the data source in this way, FusionReactor will display this name in the Data Source column of the JDBC tab in the Request Details page. This is useful if you're using multiple data sources. Below is an example:

   ```
   jdbc:fusionreactor:wrapper:{JDBC URL};name=TestName
   ```

7. In the JDBC Driver Class field, enter the name of the FusionReactor JDBC Driver Wrapper class:

   ```
   com.intergral.fusionreactor.jdbc.Wrapper
   ```

8. Enter 'FusionReactor' in the Driver Name field.
9. Select an appropriate username and password values in their respective fields.
10. Click 'Submit'. ColdFusion will immediately test the driver and provide feedback if necessary. When the FusionReactor driver is loaded, it will report “FusionReactor JDBC: Driver loaded.” to the standard output stream (usually logged to coldfusion-out.log or default-out.log within cfmx\runtime\logs)

Below are two screen shots, figure 1 covers stages 1 to 4 and figure 2 covers stages 5 to 11.

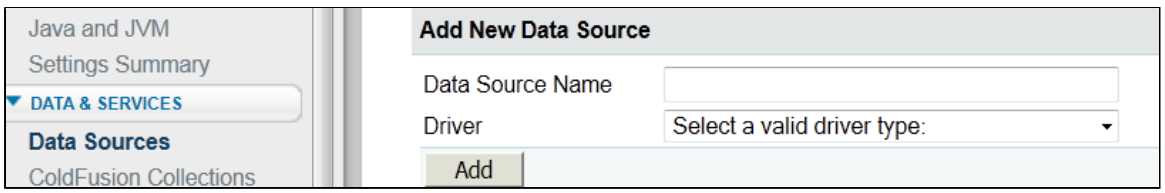


Figure 1

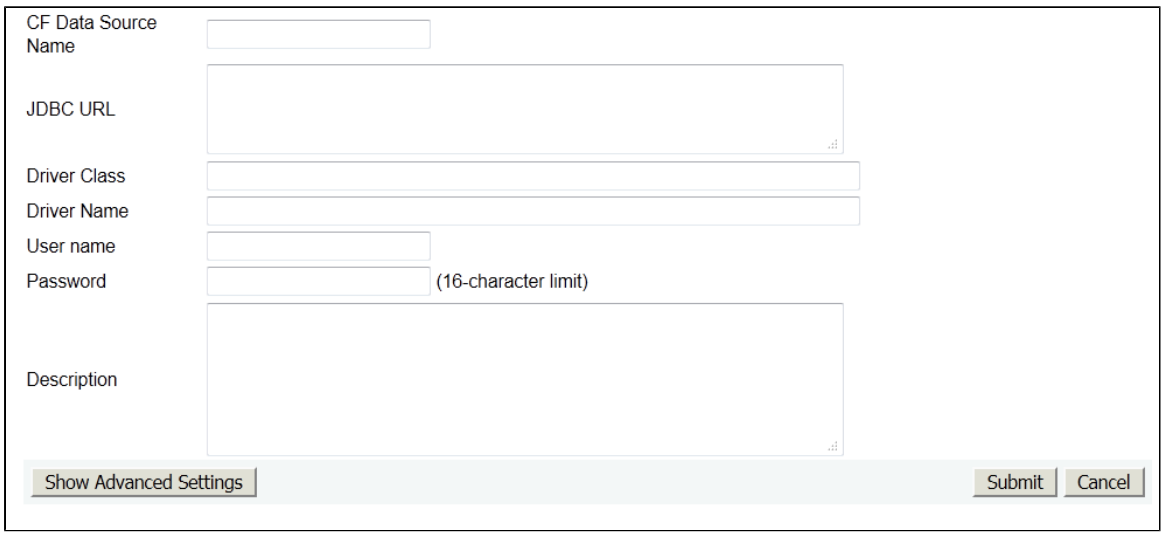


Figure 2

After submitting this form, ColdFusion will immediately test the connection and provide any necessary feedback in the ColdFusion Administrator. If the FusionReactor JDBC Driver Wrapper detects any problems with the underlying wrapped driver or with its own options, it will also provide feedback in the same way. Any exceptions generated by the FusionReactor JDBC Driver Wrapper are prefixed with an exception ID number, which may be used to look up more information. When running as part of J2EE/ColdFusion, and if the FusionReactor JDBC Driver Wrapper detects a running FusionReactor application, any queries run during the course of the request will be reported to FusionReactor for inclusion in the user interface, where appropriate.

Examples

The original data source was an MS-SQL Server with the database database name 'testdb', after being wrapped by the FusionReactor wrapped it is now called 'MyDataSource'.

Original JDBC URL:

```
jdbc:macromedia:sqlserver://localhost:1433;databaseName=testdb
```

FusionReactor URL:

```
jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://localhost:1433;databaseName=testdb};name=MyDataSource
```

Constructing JDBC URLs for other DataDirect Drivers (CFMX,7,8,9,10)

While it is not within the scope of this manual to provide exhaustive information on how ColdFusion internally constructs JDBC URLs for its built-in drivers, the following template information (taken from the neo-query.xml file within the cfmx\lib folder) may help you.

Oracle

Class macromedia.jdbc.MacromediaDriver

```
jdbc:macromedia:oracle://\[host\]:\[port\]
;SID=\[sid\];sendStringParametersAsUnicode
=\[sendStringParametersAsUnicode\]
```

MS Access

Class macromedia.jdbc.MacromediaDriver

```
jdbc:sequelink:msaccess://\[host\]:\[port\
];serverDatasource=\[datasource\]
```

PostgesQL

Class org.postgresql.Driver

```
jdbc:postgresql://\[host\]:\[port\]/\[data
base\]?\[args\]
```

MySQL

Class org.gjt.mm.mysql.Driver

```
jdbc:mysql://\[host\]:\[port\]/\[database\
]?\[args\]
```

IBM DB2

Class macromedia.jdbc.MacromediaDriver

```
jdbc:macromedia:db2://\[host\]:\[port\];DatabaseName=\[database\];sendStringParametersAsUnicode=\[sendStringParametersAsUnicode\];\[args\]
```

MS SQL Server

Class macromedia.jdbc.MacromediaDriver

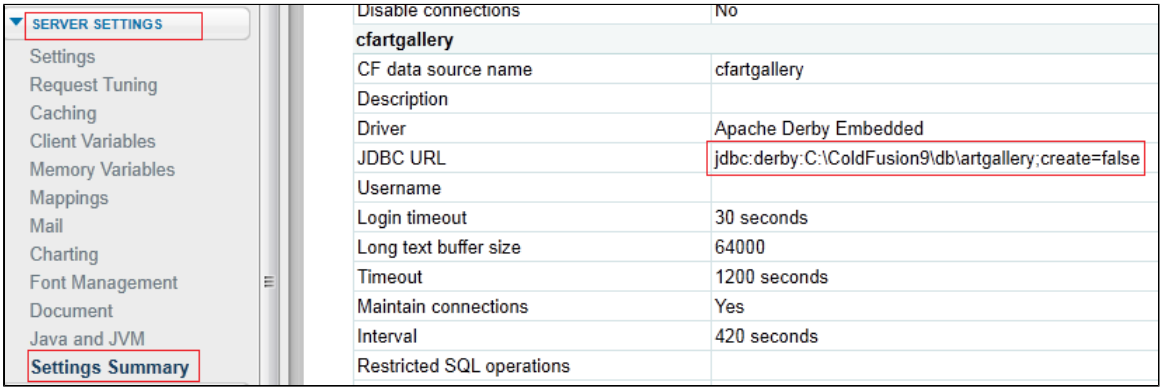
```
jdbc:macromedia:sqlserver://\[host\]:\[port\];databaseName=\[database\];SelectMethod=\[selectmethod\];sendStringParametersAsUnicode=\[sendStringParametersAsUnicode\]
```

Using a User-Specified Driver ("other") Driver

This guide walks you through locating the JDBC URL of your data source in ColdFusion and then wrapping them using the ColdFusion wrapper.

Locating your Macromedia's (DataDirect) ColdFusion Built-in Drivers

In the ColdFusion Administrator click on **Server Settings** in the table of contents, then click on the **Settings Summary** link. Locate the section with the heading Database & Services and the table with the name Database Data Sources. Locate the JDBC URL of the datasource that you want to wrap with a FusionReactor wrapper and then copy it or write it down.



Your JDBC URL will most likely be different, the purpose of the screen shot is to give you an idea of what to expect in the ColdFusion User Interface. For the purpose of this example our JDBC URL will read:

```
jdbc:sqlserver://127.0.0.1:1433;DatabaseName=testdb
```

Wrapping your JDBC URL with FusionReactor

1. Click **Data and Services** and then click on the **Data Sources** link.
2. Add a data source name.
   This page walks you through locating your Macromedia (DataDirect) ColdFusion Drivers and wrapping them in the FusionReactor wrapper. The page is broken down into 3 sections, the first walks you through finding your JDBC URL for the data source, the second walks you through wrapping it and the last one gives you an example.
3. Select 'other'-type driver.
4. Click on the Add button, you will be presented with more text input fields.
5. In the field titled CF Data Source Name enter a name for your data source.
6. In the JDBC URL text field enclosure the JDBC URL in the FusionReactor wrapped and append the JDBCs Driver Class, see below:

   ```
   jdbc:fusionreactor:wrapper:{jdbc:sqlserver://127.0.0.1:1433;DatabaseName=testdb};driver=com.microsoft.sqlserver.jdbc.SQLServerDriver
   ```

7. In the JDBC Driver Class field, enter the name of the FusionReactor JDBC Driver Wrapper class:

   ```
   com.intergral.fusionreactor.jdbc.Wrapper
   ```

8. Enter 'FusionReactor' in the Driver Name field.
9. Select an appropriate username and password values in their respective fields.
10. Click 'Submit'. ColdFusion will immediately test the driver and provide feedback if necessary. When the FusionReactor driver is loaded, it will report “FusionReactor JDBC: Driver loaded.” to the standard output stream (usually logged to coldfusion-out.log or default-out.log within cfmx\runtime\logs)

Below are two screen shots, figure 1 covers stages 1 to 4 and figure 2 covers stages 5 to 11.

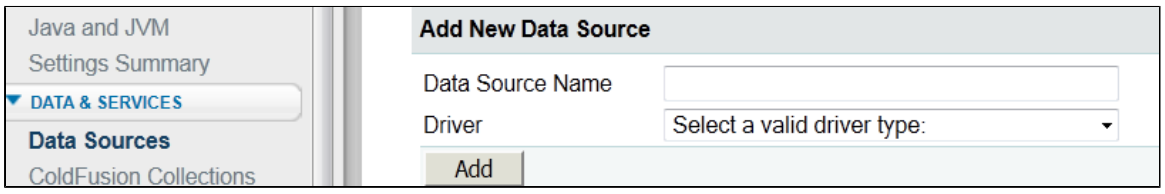


Figure 1



Figure 2

After submitting this form, ColdFusion will immediately test the connection and provide any necessary feedback in the ColdFusion Administrator. If the FusionReactor JDBC Driver Wrapper detects any problems with the underlying wrapped driver or with its own options, it will also provide feedback in the same way. Any exceptions generated by the FusionReactor JDBC Driver Wrapper are prefixed with an exception ID number, which may be used to look up more information. When running as part of J2EE/ColdFusion, and if the FusionReactor JDBC Driver Wrapper detects a running FusionReactor application, any queries run during the course of the request will be reported to FusionReactor for inclusion in the user interface, where appropriate.

Installing Third-Party JDBC JAR Files

In order for the Driver Wrapper to locate the underlying driver, the third-party JDBC JAR driver file (usually delivered as one or more JAR files) must be installed in a location visible to the Driver Wrapper classloader. We recommend either of the following locations beneath your ColdFusion installation folder, for example:

* <CFDIR>\runtime\servers\coldfusion\SERVER-INF\lib
* <CFDIR>\runtime\lib

Standalone Installation (Servlets etc.)

The driver is packaged within the fusionreactor.jar file and you must arrange for this file to be in the JVM classpath when the application runs. The URL syntax is as previously described. In this configuration, you *must* provide the 'driver' parameter, in order that the FusionReactor JDBC Driver Wrapper can load the appropriate underlying wrapped driver.

Using the FusionReactor JDBC Wrapper

Configuration Options

The FusionReactor JDBC Driver Wrapper is configured exclusively using JDBC URL driver parameters. The following parameters may be specified as pairs, separated by semicolons. No parameters are mandatory. Parameter names are case-insensitive.

driver

Value: Fully-qualified Java class name
Default: no underlying driver will be loaded.

This option instructs FusionReactor JDBC Driver Wrapper to load an underlying wrapped driver.

This is not necessary **only** if the JVM is already aware of the target driver (i.e. it has already been loaded with Class.forName("...")). Macromedia's own DataDirect drivers are loaded automatically by ColdFusion, so this option may not be necessary if you are using these drivers. However, if you are using a user-specified driver (having a JDBC URL which does not start with jdbc:macromedia), you **must** supply this parameter.

Since the registration of the driver is only ever performed once, regardless of how many connections the driver processes, this parameter can (and should) be specified on all FusionReactor wrapped Data Sources.

If you do not specify this option, and the JVM is *not* aware of the underlying driver, FusionReactor will raise an exception and the J2EE server will not verify the driver.

rowLimit

Value: integer
Default: 0 (disabled).

This option instructs the FusionReactor JDBC Driver Wrapper to limit returned rows to the given value.

After the application has retrieved this number of rows from the result set, FusionReactor will discard any remaining rows.

notifyAfter

Value: integer
Default: 0 (disabled).

This option instructs the FusionReactor JDBC Driver Wrapper to output a notification after 'n' rows have been retrieved for the query.

remindAfter

Value: integer
Default: 0 (disabled).

This option instructs the FusionReactor JDBC Driver Wrapper to periodically output a query reminder every 'n' rows. If **notifyAfter** is specified, FusionReactor JDBC Driver Wrapper will only begin reminding after the notification threshold has been reached.

E.g. notifyAfter=1000, remindAfter=100, actual rowcount 1350.
Notification occurs at row 1000, reminders at 1100, 1200 and 1300.

inhibitReformat

Value: Boolean
Default: false.

When tracking queries, the FusionReactor JDBC Driver Wrapper will reformat them for logging and presentation by attempting to make them fit on a single line.

This allows logs to be viewed more easily, but may hinder developers who are used to seeing queries formatted a certain way (as they are written in a ColdFusion page, for example). Setting this option to 'true' stops FusionReactor JDBC Driver Wrapper reformatting statement text, and allows multi-line presentation in the FusionReactor interface and log.

logToFusionReactor

Value: Boolean
Default: true.

If set to true (the default) and the FusionReactor JDBC Driver Wrapper detects a running FusionReactor instance, it will log the execution of a query to FusionReactors 'jdbc-X.log' (where 'X' is the current rolling log number).

If this option is enabled and FusionReactor was *not* detected, it has no effect.

interpretObjects

Value: Boolean
Default: true.

If set to true (the default), when a **PreparedStatement** attempts to bind an **Object** type to a positional parameter using one of the **setObject**(...) methods, the wrapper will attempt to interpret the data (for logging and reporting purposes only) by calling the **toString()** method on the object. This value will then be used in the log and FusionReactor administrator, as if the application had called a **setString**(...) method. If the object does not override the default **toString()** method, the default behavior is to return the hash code of the object.

If this parameter is false, the wrapper will use the format

```
{OBJECT java.class.name xyz}
```

where *xyz* is the .toString() representation. This makes it clear that the parameter is of type Object, but is perhaps less easy to read in the log and the Administrator.

name

Value: string
Default: empty

If specified, the driver wrapper will report metrics to FusionReactor with the given name. These names will be reported in the JDBC logfile (or as an empty value if not set). The name will also be reflected in the JDBC tab of the Request Details page, allowing the user to differentiate queries which ran against more than one datasource. This is useful when multiple databases are being used to aggregate results, or when different drivers are being tested.

cp

Value: comma-delimited list of strings pointing to jar files
Default: empty.

If specified, the driver wrapper will extend the search for the specified **driver** class to the jars pointed to by the files of this property. By default, the driver wrapper relies on the J2EE engine to locate individual driver classes, but if they are not visible to FusionReactor, you may receive the exception message **"driver class (*classname*) could not be found and loaded."** In this case, you might try adding the jar containing your driver to the **cp** property.

Users who previously used the **split-jar procedure** to work around this issue should read the section FusionReactor 4.5.0: Upgrading to 4.5.0 from the Split-Jar Procedure.

autoCommit

Value: Boolean
Default: true

If specified, the JDBC wrapper will turn off autocommit for all statements, regardless of the current status of autocommit, or the existence of any transactions. This option must be used with extreme caution as it alters the default behaviour of the JDBC system, and is provided as a workaround to J2EE servers which require it to be disabled. We do not recommend using this option to defeat autocommit. After enabling this option, you must verify the atomicity and transactional integrity of your application's JDBC statements.

How To Specify These Options

These options pertain to the FusionReactor JDBC Driver Wrapper, and should therefore be specified **outside** of the curly braces used to wrap the original JDBC URL. Any options which are required by the original JDBC URL should remain **within** the braced section.

Here's an example of a wrapped SQL Server JDBC URL, using the Macromedia driver, to which a couple of FusionReactor JDBC Driver Wrapper options have been added. The material in bold illustrates the additional wrapper syntax:

```
jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://int0007:1433;databaseName=frtest};notifyAfter=1000;remindAfter=200;inhibitReformat=true;name=DataWarehouse
```

You can see that in this example, the notifyAfter, remindAfter, inhibitReform and name options have all been specified. The databaseName option pertains to the Macromedia driver, and are therefore within the braced section.

Sample JDBC URLs

Here are a few examples of URLs, wrapped with the FusionReactor Driver Wrapper. This section is not an exhaustive reference on the syntax of each URL - you should check the documentation for each individual driver. **ColdFusion users:** in general we recommend downloading and using vendor-specific drivers if possible.

Oracle (Macromedia)

Using the Macromedia driver, with the **notifyAfter** FusionReactor Driver Wrapper option:

```
jdbc:fusionreactor:wrapper:{jdbc:macromedia:oracle://int0234:1521;SID=testdb};notifyAfter=10000
```

MySQL (Macromedia)

Using the MySQL GJT driver, with the **inhibitReformat** FusionReactor Driver Wrapper option:

```
jdbc:fusionreactor:wrapper:{jdbc:mysql://int0003:3306/webshopdb?defaultFetchSize=400};inhibitReformat=true;driver=org.gjt.mm.mysql.Driver
```

MySQL 4/5 (MySQL Commercial)

Using the MySQL commercial driver (for example, supplied by Adobe in `c:/ColdFusion10/cfusion/lib/mysql-connector-java-commercial-5.1.17-bin.jar`)

```
jdbc:fusionreactor:wrapper:{jdbc:mysql://int00d6.bbn.intergral.com:3306/TestApp?tinyInt1isBit=false&};driver=com.mysql.jdbc.Driver;name=mysqldb;cp=c:/ColdFusion10/cfusion/lib/mysql-connector-java-commercial-5.1.17-bin.jar
```

MySQL 5 Connector/J (MySQL)

Here's an example using the MySQL Connector/J driver, downloadable from mysql.com. It's installed in `c:/mysql/mysql-connector-java-5.1.20-bin.jar`.

```
jdbc:fusionreactor:wrapper:{jdbc:mysql://int00d6.bbn.intergral.com:3306/TestApp};driver=com.mysql.jdbc.Driver;name=mysqlj;cp=c:/mysql/mysql-connector-java-5.1.20-bin.jar
```

Microsoft SQL Server (Macromedia)

Using the Macromedia driver, with the **remindAfter** FusionReactor Driver Wrapper option (we've named this data source too):

```
jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://int0007:1433;DatabaseName=frtest};remindAfter=500;name=SQLServerDataSource
```

Microsoft JDBC Driver 4.0 for SQL Server (Microsoft)

Here's a URL using the Microsoft SQL Server 2005 JDBC Driver – which is not supplied by Macromedia, and so must therefore be explicitly specified as the driver together with its jar, which we've placed in `c:\jdbc`.

```
jdbc:fusionreactor:wrapper:{jdbc:sqlserver://int00d0:1433;databaseName=AdventureWorks};driver=com.microsoft.sqlserver.jdbc.SQLServerDriver;name=testdb;cp=c:/jdbc/sqljdbc4.jar
```

In this example, the username and password must be specified separately during the connection process. In ColdFusion, for example, the two fields can be entered in the DataSource Manager. Here's an example where the username and password is specified in the URL:

```
jdbc:fusionreactor:wrapper:{jdbc:sqlserver://int00d0:1433;databaseName=AdventureWorks;user=scott;password=tiger};driver=com.microsoft.sqlserver.jdbc.SQLServerDriver;name=testdb;cp=c:/jdbc/sqljdbc4.jar
```

N.b. this example won't work in ColdFusion, because the username and password must be separately specified.

You can see in these examples that we've used the **cp** option to specify the jar, where previously you would have had to use the split-jar process. Because the driver is no longer reliant on the system-supplied classloader, JDBC drivers can be loaded from anywhere, and the split-jar process is no longer required.

Derby Embedded (Apache Derby)

The Apache Derby embedded database is an in-process (no separate server) database. When this database driver is loaded, the database is started in the JVM process itself, a procedure which may only occur one time, until the database is subsequently stopped. You should therefore **only** used this database with a wrapped

datasource; you should not mix wrapped and unwrapped access to this datasource. If you *do* mix these datasources, FusionReactor will attempt to share the connection with the wrapped and unwrapped datasource - which works in most cases - though this is an **unsupported configuration**.

See Technote FRS-289: Using the JDBC Wrapper with Apache Derby Database on ColdFusion for more information.

This URL explicitly specifies the driver, together with its jar, and the **autocommit** option - also explained at the link above.

```
jdbc:fusionreactor:wrapper:{jdbc:derby:C:/ColdFusion10/cfusion/db/bookclub/;create=false;MaxPooledStatements=300};driver=org.apache.derby.jdbc.EmbeddedDriver;name=bookclubdb;cp=c:/ColdFusion10/cfusion/lib/derby.jar;autocommit=false
```

And Finally...

... and as a quick reminder, the FusionReactor Driver Wrapper driver class is:

com.intergral.fusionreactor.jdbc.Wrapper

Interpreting JDBC Log Data

When outputting data to the JDBC log file (which can be found in FusionReactor's instance log directory), the FusionReactor JDBC Driver Wrapper outputs a number of fields which can be used to debug JDBC transactions and derive statistics about how the system is using database resources.

The JDBC log file is space-delimited, with text fields (which may contain spaces) enclosed with double-quotes. We have had no trouble importing this data into Microsoft Excel and OpenOffice Calc, as well as Microsoft SQL Server using Data Transformation Packages.

The following list describes the meaning of each field. The list describes the fields in left-to-right order. For field sources listed as 'FusionReactor', this field may be empty **if** the request in which the query ran has no associated FusionReactor tracked request.

Calendar Date

Value: YYYY-MM-DD
Source: Wrapper

Specifies the calendar date on which the log message was raised.

Time

Value: HH:MM:SS
Source: Wrapper

Specifies the 24-hour time at which the log message was raised.

Epoch time

Value: long millisecond
Source: Wrapper

Specifies the exact epoch time (millisecond offset from midnight on January 1st 1970 UTC) at which the log message was raised.

Fusion Request ID

Value: long integer
Source: FusionReactor

Specifies the FusionReactor request ID within whose execution this JDBC interaction occurred.

Thread

Value: String
Source: Wrapper

Specifies the name of the thread in which this JDBC interaction occurred.

Client IP

Value: dotted quad IP address
Source: FusionReactor

Specifies the IP of the client for whom this request is running.

HTTP Method

Value: HTTP 1.X Method (GET / POST / HEAD etc.)
Source: FusionReactor

Specifies the HTTP method of the request which caused this JDBC interaction

URL

Value: Full or Partial URL
Source: FusionReactor

Specifies the URL which caused this request to run. If FusionReactor is tracking complete URLs (which machine name) this will be a full URL. If not, this will be the path element.

Log Message Type

Value: One of METRIC, NOTIFICATION or REMINDER
Source: Wrapper

Specifies the type of this message. *METRIC* reports the completion of a JDBC interaction, NOTIFICATION specifies that a Notification threshold has been reached on the size of the result set, and REMINDER specifies that a result set size reminder interval has been reached.

Execution Start Time

Value: long millisecond
Source: Wrapper

The start time in milliseconds from the epoch datum (see **Epoch Time** above) at which the JDBC interaction began (i.e. the time at which the statement was transferred to the database driver for execution)

Execution End Time

Value: long millisecond
Source: Wrapper

The time at which the underlying database driver finished executing the statement.

Result Set Close Time

Value: long millisecond
Source: Wrapper

The time at which the result set was closed by the J2EE application (e.g. ColdFusion etc.). This interval between this time and the Execution Start Time is useful as the total processing time for the query, including database execution time and the time taken for the J2EE application to fully read and process the result set.

Execution Elapsed Time

Value: long millisecond
Source: Wrapper

The time taken to execute the statement on the database (computed from Execution Start and End times)

Result Set Elapsed Time

Value: long millisecond
Source: Wrapper

The time taken between sending the statement to the underlying driver for execution, and the J2EE application actually closing the result set (computed from the Execution Start time and the Result Set Close Time)

Rows Read

Value: long
Source: Wrapper

Specifies the maximum number of rows read by the J2EE application. If the statement is not a DQL command (select etc.), but rather is DML/DDL/RIGHTS (insert/update, drop/alter/create, revoke/grant etc.) this value will be 0.

Is Prepared Statement

Value: boolean
Source: Wrapper

Specifies whether this statement was prepared in advance of its execution.

Is Row Limited

Value: boolean
Source: Wrapper

Specifies whether the Row Limiter activated to stop the query.

Datasource Name

Value: String
Source: Wrapper

Specifies the name of the datasource (specified by the JDBC 'name' parameter). Blank if the name was not specified.

Statement

Value: SQL String
Source: Wrapper

Specifies the statement which was run during this interaction. Any whitespace formatting in the original statement is flattened to allow the statement to appear on one line. If the interaction was a batch execution, individual statements are delimited by [[and]] strings.

Stack Elements

Value: Comma-separated list of Strings
Source: Wrapper

If FusionReactor is configured to record stack traces, this field contains a comma-separated list of stack frames, recorded when the query completed. If the debug information is available, this field can be used to locate the exact line in a script or Java program which caused the interaction.

URL Parameters

Value: String
Source: FusionReactor

Contains the parameters which were present on the URL associated with the request in which this statement is running.

Message

Value: String
Source: Wrapper

For NOTIFICATION or REMINDER log messages, this field contains the text of the notification or reminder.

Prepared Statement Positional Bind Parameters Replacement Strings

During presentation of the statement text for Prepared Statements, the FusionReactor JDBC Driver Wrapper attempts to replace the bind parameter placeholder (“?”) with an appropriate representation of the bound value for that parameter. This is not always possible, for example with binary streams. The following table details the replacement.

Textual replacements in italics have no feasible representation and are inserted as literals. All other types are converted to their string representation for output.

Parameter Type	PreparedStatement Method	Text Representation
SQL Array	setArray(...)	*{SQL ARRAY}*
ASCII Stream	setAsciiStream(...)	*{ASCII STREAM}*
Binary Stream	setBinaryStream(...)	*{BINARY STREAM}*

Big Decimal	setBigDecimal(...)	String representation
Blob	setBlob(...)	*{BLOB}*
Boolean	setBoolean(...)	String representation
Byte	setByte(...)	String representation
Bytes	setBytes(...)	String representation: comma-separated list of up to 256 bytes
CharacterStream	setCharacterStream(...)	*{CHARACTER STREAM}*
Clob	setClob(...)	*{CLOB}*
Date	setDate(...)	String represenation using Date.toString() to format
Double	setDouble(...)	String representation
Floating Point Number	setFloat(...)	String representation
Integral Number	setInt(...)	String representation
Long Integral Number	setLong(...)	String representation
Null	setNull(...)	*{NULL}*
Object	setObject(...)	String representation **if** option **interpretObjects** is true (or not specified) **otherwise:** “{OBJECT java.class.name xyz}” where *xyz* is the .toString() representation
Reference	setRef(...)	*{REF}*
Short Integral Number	setShort(...)	String representation
String	setString(...)	String representation
Time	setTime(...)	String representation
Timestamp	setTimestamp(...)	String representation
Unicode Stream	setUnicodeStream(...)	*{UNICODE STREAM}*
Uniform Resource Locator	setURL(...)	String representation

New in FusionReactor 4.5.1

Parameter Type	PreparedStatement Method	Text Representation
Input Stream (Binary)	setBinaryStream(...)	*{BINARY STREAM}*
Input Stream (BLOB)	setBlob(...)	*{BLOB STREAM}*
Reader (Character)	setCharacterStream(...)	*{CHARACTER STREAM}*
Reader (CLOB)	setClob(...)	*{CLOB STREAM}*
Reader (NChar)	setNCharacterStream(...)	{NCHAR STREAM}
NClob	setNClob(...)	*{NCLOB}*
NString	setNString(...)	*{NSTRING}*
RowID	setRowId(...)	*{ROWID* x} "x" is the .toString() representation of the Row ID
SQLXML	setSQLXML(...)	*{SQLXML x}* "x" is the .toString() representation of the SQL XML

A Note On SQL Server Select Methods

In this edition of the JDBC Driver Wrapper manual, we have removed references to the Macromedia/Microsoft SQL Server driver option **selectMethod**. This option controls how the driver and the database supply rows to your application, and can have an impact on memory usage and JDBC processing speed. We want you to make an informed decision about whether to apply this option, and this section will provide an explanation about how this option works.

When You Can Use This Option

This option is applied as a driver parameter in the JDBC URL for **Microsoft SQL Server datasources**. It can be used with the Microsoft JDBC Driver (2000 and 2005), and the built-in Macromedia driver. Since it is a parameter to the underlying JDBC driver, it must be supplied within the wrapped part of the URL:

Using the Macromedia Driver:

```
jdbc:fusionreactor:wrapper:{jdbc:macromedia:sqlserver://int0007:1433;DatabaseName=frtest;SelectMethod=cursor};remindAfter=500;name=SQLServerDataSource
```

Using the Microsoft SQL Server 2005 JDBC Driver:

```
jdbc:fusionreactor:wrapper:{jdbc:sqlserver://int0007:1433;DatabaseName=frtest;SelectMethod=cursor};driver=com.microsoft.jdbc.sqlserver.SQLServerDriver
```

Direct and Cursor Selection Modes

If this option is not specified in the JDBC URL, it defaults to **direct**.

Direct Selection Method

When the driver is operating in direct mode, select statements are processed like this:

- Application performs an SQL select.
- JDBC driver transfers the request to the database.
- Database performs the select.
- **Database transfers the complete result set back to the driver**.
- Driver makes each row available to the application.

Cursor Selection Method

Here is the same process when the driver operates in cursor selection mode:

- Application performs an SQL select.
- JDBC driver transfers the request to the database.
- Database performs the select.
- **Database opens a cursor, and sends the first batch of rows back**.
- Driver makes the first batch of rows available.
- Application eventually requests a row which wasn't in the first batch.
- Driver asks the database for more rows.
- **Database uses the cursor to see where it was in the results, and transfers the next batch.**

.. etc., until the application closes the query or the database exhausts the cursor.

Pros and Cons

Direct mode transfers all the data at once. The data is then stored in the driver's memory space.

Pros

- Very fast for the application to iterate through rows because all data is immediately available.
- Useful for fast batch work where speed is preferable over low memory usage.

Cons

- Large result sets will consume a lot of memory. You must account for this memory when configuring your J2EE (ColdFusion) application server. You must provide enough margin for the size of the largest result set to be requested.
    - Repeatedly opening large result sets may cause Java to perform major garbage collections often, causing the application to become less responsive.
    - Concurrently opening large result sets may cause memory to be exhausted, leading to continuous garbage collection as Java attempts to reclaim memory, and possible crashes.
    - Be aware that the total memory requirement for a given result set is **twice the size of the result**

set. When the database runs the query, it must first buffer the results in its own memory space, before transferring all the data to the client.

- Long network delay when query is run while all data is transferred. For databases running on the same server, this delay is shorter but not insignificant, since the TCP/IP transport is still used.
- Large time and memory overhead if not all rows are processed by the application.
    - If the application only processes the first few rows, the transfer and storage of the entire result set is very inefficient.
- Concurrent requests for result sets from the database will cause separate connections to be opened to the database.

Cursor mode transfers data in batches transparently to the application. The complete resulting row set for each query is stored in the database. The J2EE (ColdFusion) application which uses the JDBC driver only stores one batch (typically a small fraction) of rows. The driver takes care of requesting new batches from the database as they are needed.

Pros

- Small memory requirement in the client application, since only one batch of rows is stored.
    - Memory demand burden is placed on the database, which can typically manage it more efficiently than J2EE applications. If the database is running on a dedicated machine, the memory requirement is then transferred to that machine, leaving the J2EE (ColdFusion) server with more memory to allocate to the application.
- Result set processing can begin much more quickly since there is a much smaller initial delay before the driver makes rows available.
- Repeatedly opening large queries, running large concurrent queries, or only processing the first few rows of a large result set is comparatively cheap, since only the first batch will be transferred and stored in the client.

Concurrently-running queries may be interleaved over the same connection.

Cons

- Total time taken to retrieve an entire result set is slightly higher than in direct mode, since the driver must retrieve several batches of rows.

Caveats for non-ColdFusion JDBC Environments

This section is for engineers using SQL Server JDBC drivers in a J2SE/J2EE environment which is not ColdFusion; e.g. where the driver, connection etc. is obtained manually using Java code.

This caveat applies only to MS SQL Server 2000 JDBC Driver and DataDirect/Macromedia SQL Server Drivers.

This problem does not occur with the MS SQL Server 2005 JDBC Driver.

Connection Commit Status Restrictions in Direct Mode

You may see the following JDBC exception:

> **“Can't start a cloned connection while in manual transaction mode”.**

When operating in direct mode, only one active SQL statement (including 'select' statements) can be open over a *java.sql.Connection*.

If a second or subsequent statement is opened while any result sets or statements are still active on the first connection, the connection is transparently cloned by the driver, and the statement is run over the new connection.

If the auto-commit status of the connection has been changed from the default *auto-commit=true* to a*uto-commit*

=*false* (that is, all transactions will be explicitly committed or rolled back), the cloned connection may have uncommitted (dirty) data present. In this case it is not safe to start a new statement, and the driver issues the exception above.

To alleviate this situation, one or more the following measures may be taken:

1. Switch to cursor mode by appending *selectMethod=cursor* to your JDBC URL.
2. Switch to auto-commit mode.
3. Ensure that only one statement is active on a connection by closing all statements and result sets prior to opening new ones.
4. Ensure that all connection operations are synchronized against multiple access in a multi-threaded environment.

Furthermore, several operations may cause the driver to internally create a second statement and attempt to run it on the same connection (e.g. updating certain types of result set), producing the same exception.

References

Microsoft KB article 313181 refers: http://support.microsoft.com/kb/313181

Exception Catalog

In almost all cases, the FusionReactor JDBC Driver Wrapper will propagate all SQL Exceptions upwards, allowing them to be transparently handled by an application. In certain cases, FusionReactor JDBC Driver Wrapper can originate its own exceptions.

These exceptions have an ID number placed in square brackets within the message string. This ID number can be used to locate more information in the following table.

ID	Example	Explanation
1	Driver options must be pairs	Indicates an error with the wrapper driver parameters. Check all options have values.
2	value for driver option '(*option*)' was neither 'true' nor 'false' (*value*)	The valid values for this option are either 'true' or 'false'. The actual read value is provided in the message. Check the option specified has a valid value.
4	driver option '(*option*)' is unknown	The supplied option was not known. The rogue option is given in the message. Check the option for typographical errors.
5	couldn't parse value for driver option (*option) (value)* as a number.	The given driver option requires a number as its value, and FusionReactor couldn't parse the given value as a number. The option and the rogue value are given in the message. Check the value to make sure it can be parsed as a number and has no alphabetic characters.

6	driver class (*classname*) could not be found and loaded.	The Java class specified with the 'driver' option could not be found. Check the class or jar is within the application classpath, and the class name is fully qualified (if required). If you wish to extend the classpath, use the **cp** driver option to specify further jars (see Using the FusionReactor JDBC Wrapper for more information on the **cp** opti on.)
7	Could not find a wrapped JDBC URL within the passed string. The wrapped URL must be within braces { }.	The Driver Wrapper could not find a wrapped URL within the JDBC URL. Check the syntax of the whole URL and make sure the original URL is specified within braces ....
8	Soft kill active for this thread	This thread has been marked for Soft Kill by FusionReactor. The Driver Wrapper will not proceed with any further database activity.
9	URL supplied did not conform to JDBC URL specification. Please check the syntax.	The Driver Wrapper could not reliably decipher the supplied JDBC URL, because it probably contained a syntax or typographical error. Please check it carefully by hand.

FusionReactor 4.5.0: Upgrading to 4.5.0 from the Split-Jar Procedure

Background

Because of the way J2EE classloaders are partitioned, in some configurations, FusionReactor's JDBC Wrapper was unable to wrap a driver because the classloader or JDBC Driver Manager forbade access to it.

The solution in all versions of FusionReactor **prior** to 4.5.0 is to apply Technote FRS-240, "Standalone FR JDBC Wrapper Driver (for MySQL Connector/J etc.)", commonly known as the **split-jar procedure**. This procedure split the FusionReactor system into two separate jars. One jar, fusionreactor-jdbc.jar was placed with the underlying JDBC driver jar. The second jar, fusionreactor-core.jar replaced the monolithic fusionreactor.jar, which was deleted. Because the -jdbc component could locate the underlying driver, and the two components could communicate, JDBC wrapper functionality was restored.

FusionReactor 4.5.0

FusionReactor 4.5.0 introduces a new option to the JDBC driver, cp, which is a comma-separated list of file paths pointing to jar files which contain the underlying driver to be wrapped.

It is no longer necessary to split the jar, and the FusionReactor 4.5.0 instance updater will in fact restore the monolithic FusionReactor jar file `fusionreactor.jar`, and **remove** the core and JDBC split jars from the system. You must then modify your datasource to add the **cp** option to point to the actual JAR which contains

the driver you're wrapping.

ColdFusion users, for example, can do this using the **Data & Services -> Data Sources** tool within the CF Administrator. Many examples are provided in the JDBC Wrapper User Manual, see here for some examples.

For instance, if you were previously using a wrapped Microsoft SQL Server Driver, you must append a CP option which points to the jar. In this example, the jar has been installed to `c:/jdbc/sqljdbc4.jar`:

```
jdbc:fusionreactor:wrapper:{jdbc:sqlserver://int00d0:1433;databaseName=AdventureWorks};driver=com.microsoft.sqlserver.jdbc.SQLServerDriver;name=testdb;cp=c:/jdbc/sqljdbc4.jar
```

Users of the Apache Derby database on ColdFusion must also use the **`autoCommit=false`** option to prevent CF committing the transaction before all data has been read. FRS-289, "*Using the JDBC Wrapper with Apache Derby Database on ColdFusion*" refers.

FusionReactor Plugins

FusionReactor 4 introduces a new OSGi based plugin architecture to enable new features to be seamlessly deployed and ensuring we can enhance the monitor with additional functionality that you request in the future.

AMF Dissector Plugin

FusionReactor Deadlock Plugin

FusionReactor Extensions for ColdFusion (FREC)

- FREC Installation Guide
  - Installing FusionReactor Extensions for ColdFusion (FREC)
  - Installing FREC into FusionReactor Instances
  - Uninstalling FREC from FusionReactor Instances
  - Uninstalling FusionReactor Extensions for ColdFusion (FREC)
- FREC Log Files
  - CF Database Connection Statistics (dbpoolstats.log)
  - CF Hit Count Statistics (hitcountstats.log)
  - CF Log Entries (fr-coldfusion-log.log)
  - CF Realtime and Request Stats (realtimestats.log)
- Configure FusionReactor ColdFusion Plugin (FREC)
- How to enable FREC

FusionReactor Heartbeat Plugin

FusionReactor JRun Plugin

FusionReactor Log Rotator Plugin

FusionReactor Notification Generator Plugin

AMF Dissector Plugin

Description

The Action Message Format (AMF) Dissector provides a means to see the AMF messages being sent to and from the server from within FusionReactor's Request Detail page. AMF is a binary format used to serialize ActionScript objects. It is used primarily to exchange data between an Adobe Flash application and a remote service, usually over the internet. The AMF tab can be found under the Request Detail Page.

FusionReactor makes it easy by decoding your streaming AMF connections to help you debug your code.

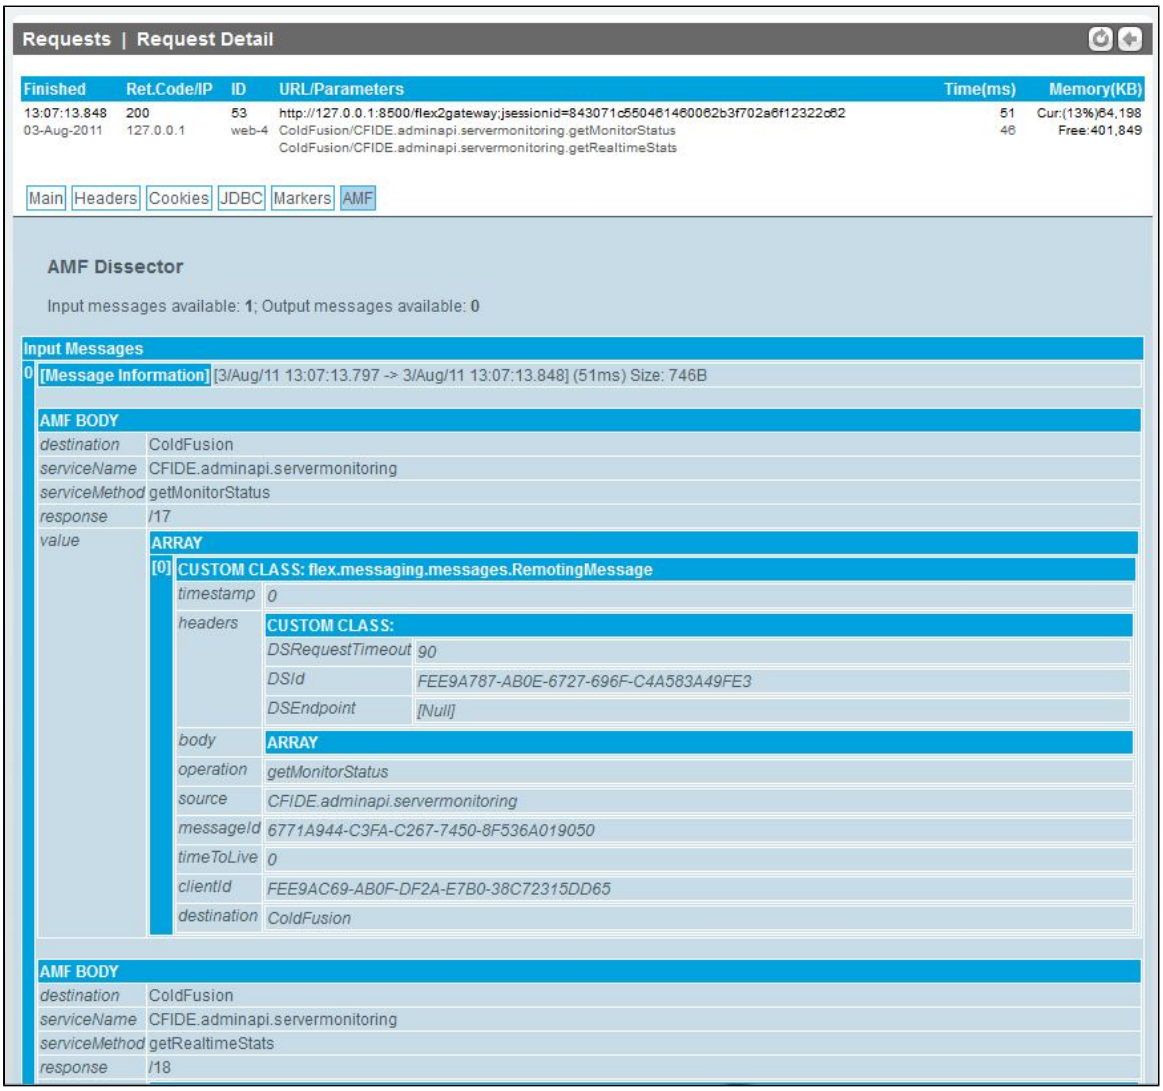


Figure 1: Decoded AMF details.

Configuration

The AMF Dissector Plugin | Configuration (accessible from the Active Plugins page) allows you to alter the following:

Field Name	Description

Dissector Mode	This allows you to record either Input AMF, Output AMF, Both or None. Decoded AMF data can be viewed by going on Request History and clicking the detail button ( ) and then clicking on the tab that is named "AMF" (as seen in Figure 1).
Buffer Size (bytes)	This is the buffer size you can allocate to each message. It is advisable that you select a reasonable size otherwise data may get rejected by the dissector.
Number of Messages	This is the number of messages in an AMF stream that can be stored per transaction.
Streaming Message Store Mode	**Store First Messages** When the buffer is full the most recent messages will be discarded. **Store Latest Messages** When the buffer is full the older messages will be discarded.
Message Display Ordering	This allows you too choose the order that messages are displayed, choosing Newest First means that the most recent messages are displayed first. Choosing Oldest First means that the oldest messages will be displayed first.

Log File

No log files about captured data are generated from this plugin.

Status Log - plugin-fr-amf-decorator-plugin.log

This plugin creates a log file called **plugin-fr-amf-decorator-plugin.log** which contains information on the status of the plugin, when it started, stopped, if there were any problems. Below is a table explaining what all the attributes in the log file mean.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.

Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information, where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as that logging has started.

Also See

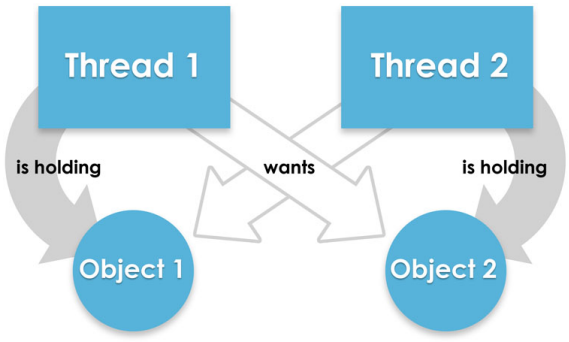
Log Settings
FusionReactor Plugins

FusionReactor Deadlock Plugin

Description

The FusionReactor Deadlock plugin polls all Java threads running in the JVM that your server is running on every 30 seconds. Any Java thread deadlocks detected are written to a log file (deadlock.log, see below) and a notification (see below) is created for that particular deadlock.

A deadlock occurs when two or more threads are waiting for each other to finish. For example Thread-1 may have a hold on Object1, Thread-2 may have a hold on Object2. In order for Thread-1 to complete it's operation it needs to access Object2, for Thread-2 to complete it's operation it needs to access Object1. Both threads are in a deadlock as they cannot proceed until they get the resource they require, and they cannot as the other Thread won't give up that resource until they have completed their operation. Thread-1 may even have a hold on other resources (for example Object3, Object4 and Object5) that other threads require causing further deadlocks.



The diagram accompanies the text above.

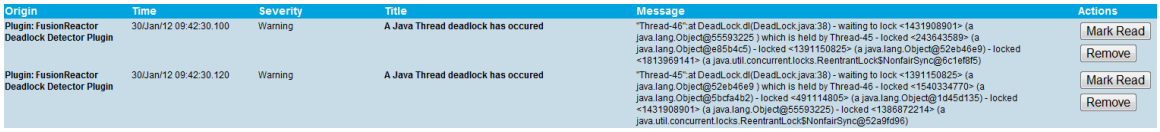
There is no specific effect that a deadlock will have on your application as there are no two identical applications. Some of the effects of deadlocks can include slow performance and high memory consumption, in sever cases they can bring your application to a complete halt. Reading the notifications this plugin generates can help you pinpoint where deadlocks are occurring in your application, with this knowledge you can make a decision as to how to proceed.

NOTE

To use this plugin you need to be using Java 1.5 or greater and FusionReactor 4.0.6 or above

Notification

If the plugin detects a deadlock a notification will be created. The notification contains information about the deadlock, this information can be used to determine how the deadlock occurred.



The image shows what a typical deadlock notification looks like

Log File - deadlock.log

When a deadlock occurs the plugin generates a file called deadlock.log and populates it with information which can be used to determine how the deadlock occurred.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date the log was written.
Time (Formatted)	2 [B]	The time the log was written.
Date/Time (ms)	3 [C]	The number of seconds since epoch.
Version	4 [D]	Version number of this log entry.
Server Start up Time (ms)	5 [E]	Start up time of the application server in milliseconds (ms) since epoch.
Log Type	6 [F]	This will either be "DEADLOCK" or "DEADLOCKRESOLVED", it describes the nature of the entry.
Thread Name	7 [G]	The thread name of the thread that has entered deadlock.
Thread ID	8 [H]	The thread ID of the thread that has entered deadlock.
Thread State	9 [I]	The state of the thread that has entered deadlock.

Lock Identity Hash Code	10 [J]	The identity hash code of the lock object. **For this field to be written to the log file you need to be using Java 1.6 or greater.**
Lock Name	11 [K]	The name of the lock.
Lock Owner Name	12 [L]	The name of the lock owner.
Lock Owner ID	13 [M]	The ID of the lock owner.
Is Native	14 [N]	If the thread that is in deadlock is running native code then this value is set to true, otherwise it is set to false.
Is Suspended	15 [O]	If the thread that is in deadlock is suspended then this is set to true, otherwise it is set to false.
Stack Trace	16 [P]	This is a full stack trace of the thread that is in deadlock.
Threads Locked Monitors	17 [Q]	This is a list of objects that are currently under lock by the thread. Each entry in the list is separated by a comma . The format for each entry is *identity hash code: string representation of the lock,* for example "675587393:java.lang.Object@2844a541". **For this field to be written to the log file you need to be using Java 1.6 or greater.**
Threads Locked Synchronizers	18[R]	This is a list of locks that are currently under lock by the thread. Each entry in the list is separated by a comma . The format for each entry is *identity hash code: string representation of the lock,* for example "675587393:1475122501:java.util.concurrent.locks.ReentrantLock$NonfairSync@57ec9545". **For this field to be written to the log file you need to be using Java 1.6 or greater.**

Status Log - plugin-fr-deadlock-plugin.log

The Deadlock plugin creates a log file called plugin-fr-deadlock-plugin.log which contains information on the status of the plugin, when it started, stopped and if there were any problems. Below is a table explaining what all the attributes in the log file mean.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information, where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.

Also See

Notifications
Log Settings
FusionReactor Plugins

FusionReactor Extensions for ColdFusion (FREC)

Introducing FREC

FusionReactor Extensions for ColdFusion (FREC) is a Plugin for FusionReactor. FREC works with ColdFusion Standard AND Enterprise (ver. 7, 8, 9 and 10) and extends FusionReactor to enable capture of the CF event, server and application logs, as well as real-time stats. The real-time stats consist of 40 new metrics including Active Sessions, # of current sessions, App Scope Size, hit counts (breakdown of hits per request type) and DB Pool Stats (detailed metrics for DSN's) which increase the amount of information that FusionReactor generates when dealing specifically with ColdFusion servers. This information is delivered to you via log files which can be used to create detailed analysis of your server by using software such as FusionAnalytics. In order to use this plugin you will have to first install the "FusionReactor ColdFusion Extensions Instance Installer". Once that is installed you need to run it and it will allow you to install the FREC plugin into any available FusionReactor Instances.

The following log files are captured by FREC:

- dbpoolstats.log
- hitcountstats.log
- realtimestats.log
- plugin-fr-logrotator-plugin.log

Downloading FREC

You can download FREC for free by visiting the plugins page on the FusionReactor website. Clicking this link will take you there.

Installing FREC

In order to install FREC, you first need to download it. Make sure to download the correct version for your operating system.
Once you've downloaded the installer, run it and follow the installation page for a more detailed guide on the installation.

When you have installed FREC, you will want to install it into an instance of FusionReactor. You will need to run the Manage Instances application which can be found in the table below (if you installed it to the default directory), depending on your operating system:

Platform	Default Location
Windows	`C:\FusionReactorColdFusionExtensions`
Linux	`/opt/FusionReactorColdFusionExtensions`
Solaris	`/opt/FusionReactorColdFusionExtensions`
Mac OS	`/Applications/FRECF`

Please see Installing FREC into FusionReactor Instances for a detailed guide on how to install FREC into a FusionReactor Instance. Once installed there are a few simple things that need to be enabled in order for FREC to work, please see How to enable FREC. This page goes through the things that must be enabled.

Uninstalling FREC

In order to uninstall the FusionReactor Extensions for ColdFusion (FREC) plugin from an instance of FusionReactor you need to run the FREC Instance Installer which can be found in the table below, depending on your operating system.

Once you run the Manage Instance application, you will need to change the action to uninstall and click "next" to start the uninstall.
For a more detailed guide, please see Uninstalling FREC from FusionReactor Instances

To uninstall the manage instances application, you will need to navigate to the directory in the table below (if you installed it to the default directory), depending on your operating system.
Run the uninstaller. For a detailed guide on how to uninstall FREC completely, please see Uninstalling FREC

> ⛔ Uninstalling the manage instance application will uninstall everything to do with FREC.

Platform	Default Location
Windows	`C:\FusionReactorColdFusionExtensions`
Linux	`/opt/FusionReactorColdFusionExtensions`
Solaris	`/opt/FusionReactorColdFusionExtensions`
Mac OS	`/Applications/FRECF`

FREC Installation Guide

Installing FusionReactor Extensions for ColdFusion (FREC)

Installing FREC into FusionReactor Instances

Uninstalling FREC from FusionReactor Instances

Uninstalling FusionReactor Extensions for ColdFusion (FREC)

Installing FusionReactor Extensions for ColdFusion (FREC)

In order to install the FusionReactor Extensions for ColdFusion (FREC) plugin into an instance of FusionReactor you will need to first need to install the FREC Instance Installer. In order to do this, run the installer (eg. `FusionReactorColdFusionExtensionsSetup_windows_1_0_0.exe` on Windows).



Figure 1: FREC Welcome Screen

The first screen in the installer is just the Welcome screen. Click the '*Next* >' button to continue on or click on '*Ca ncel*' to stop the installation.

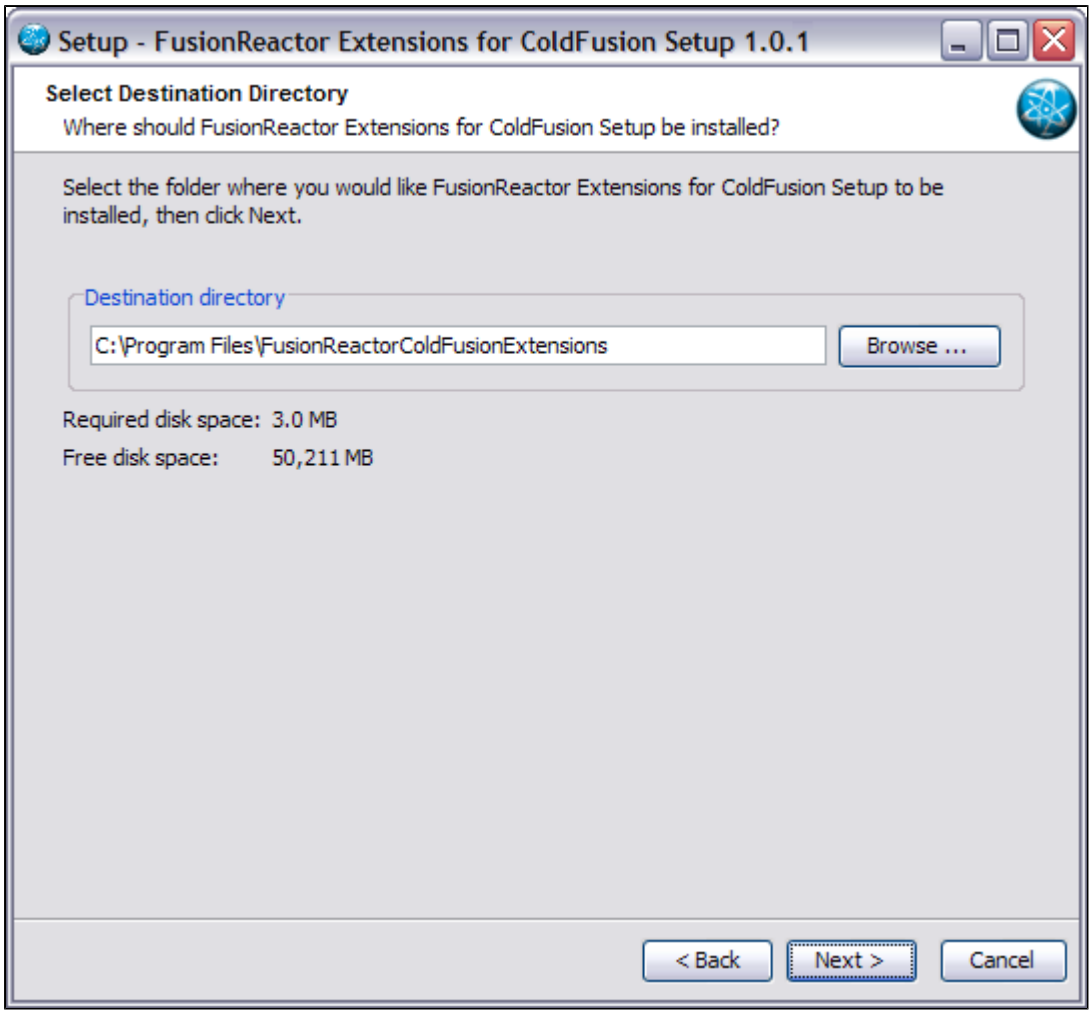


Figure 2: Destination Directory Choice

The second screen of the installer prompts you to select a destination directory for the Instance Installer. A default directory will be preselected unless you modify the value if you wish. Click the '*Next* >' button to continue, click the '< *Back*' button to back to the previous screen or click '*Cancel*' to stop the installation.

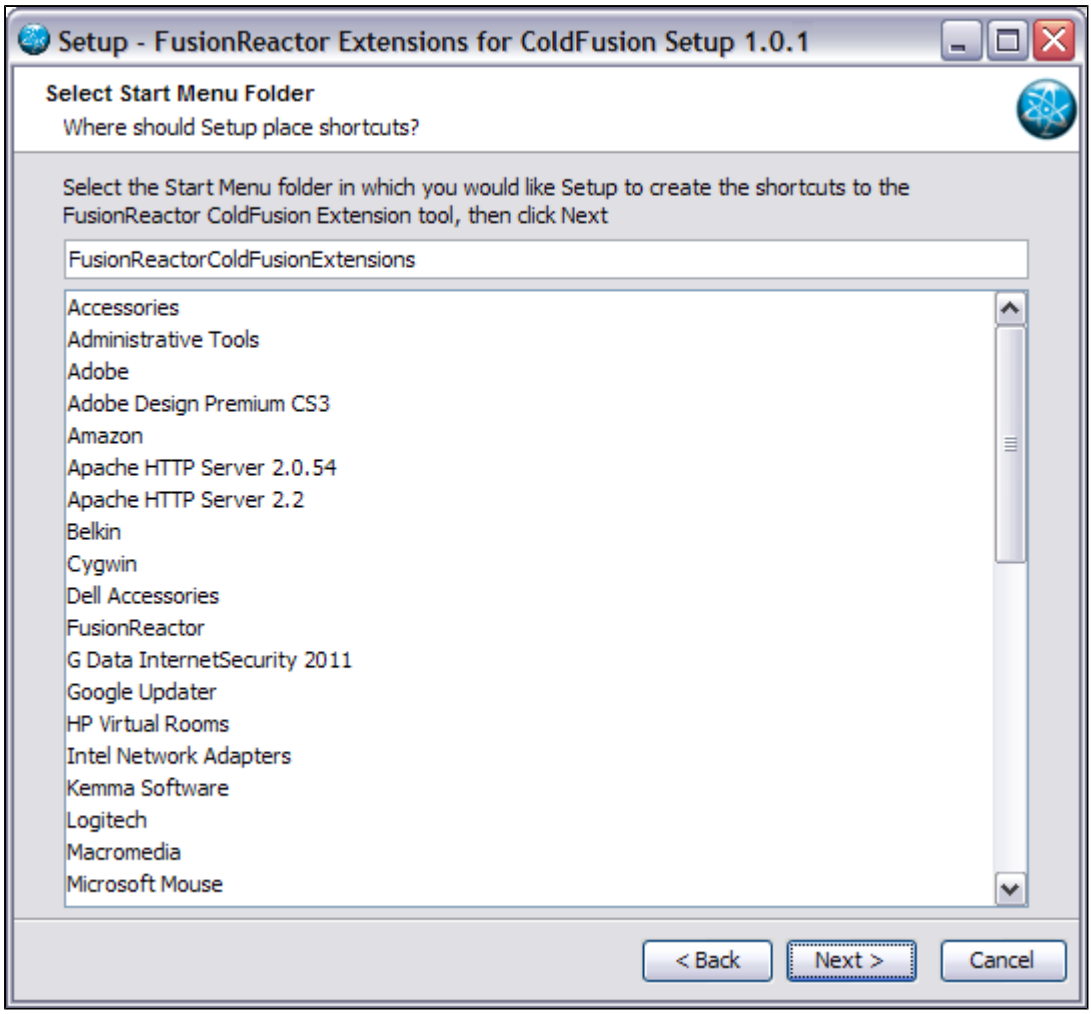


Figure 3: Select Start Menu Folder

In the third screen, you are prompted to select the Start Menu folder. Just click '*Next >*' if you want to use the default values, '*< Back*' to go back to figure 2 or '*Cancel*' to stop the installation.

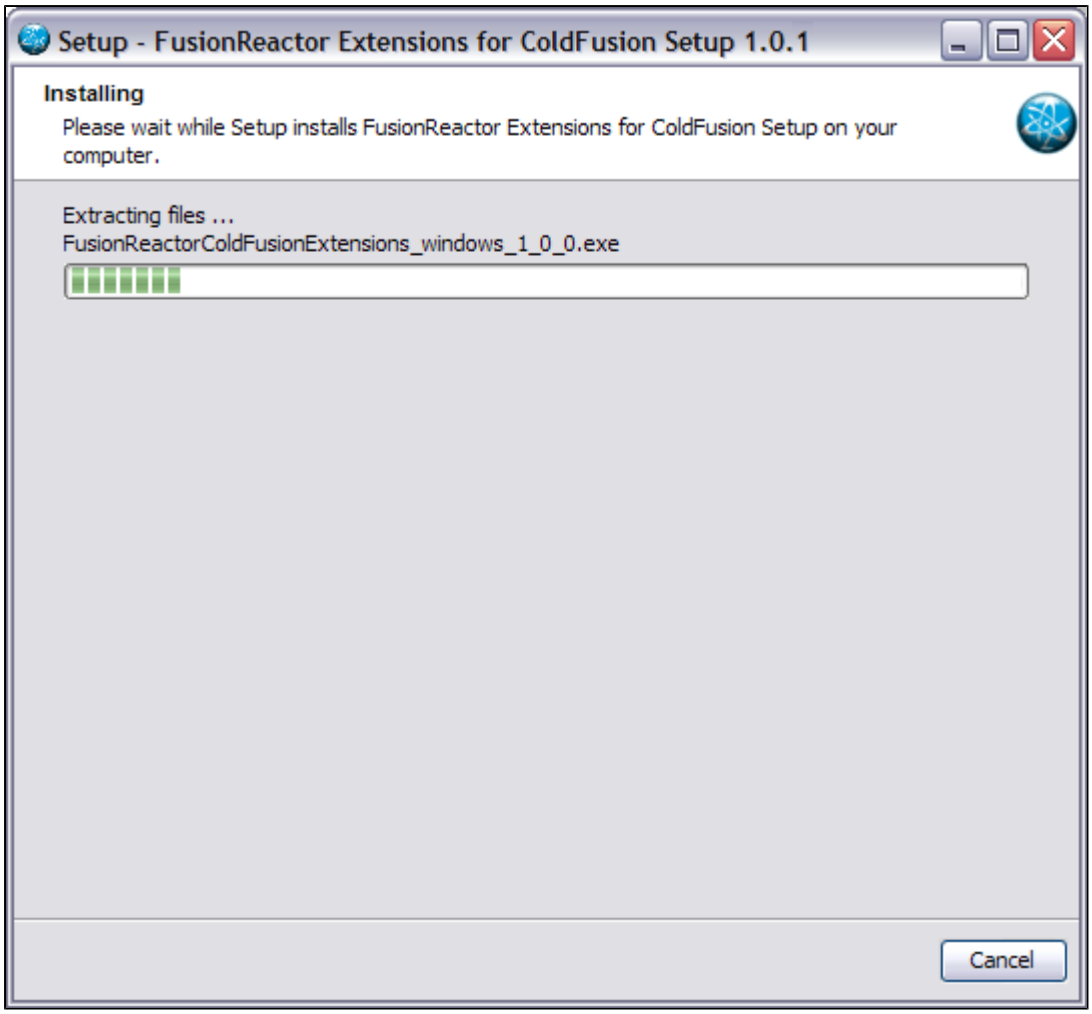


Figure 4: Installation

Once you have selected the values from the previous installer screens you will see the Progress page (Figure 4).

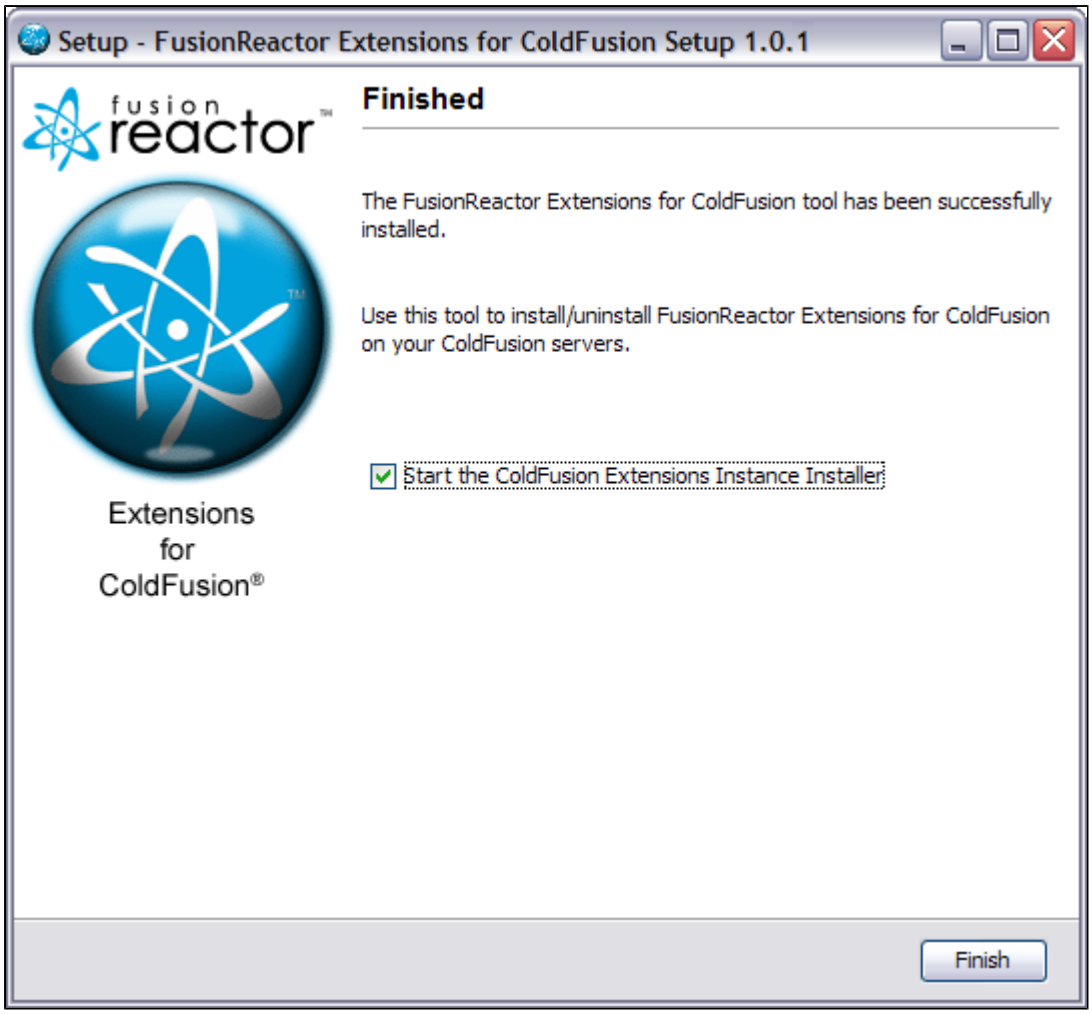


Figure 5: Finished Screen

Once the installation is completed you will see the final screen. By default you will be given the option to Start the ColdFusion Extensions Instance Installer as soon as the Installation is completed. Uncheck this box if you don't want to run the application that installs FREC into coldfusion servers. Click the '*Finish*' button to complete the installation and start installing FREC into FusionReactor instances.

Installing FREC into FusionReactor Instances

Once you have installed the Manage Instances application, run the manage instances application to install the actual plugin into your ColdFusion Instances in which FusionReactor has been installed. Once you have run it (eg. FusionReactorColdFusionExtensions -> Manage from the Windows Start Menu) you will see the following screen:

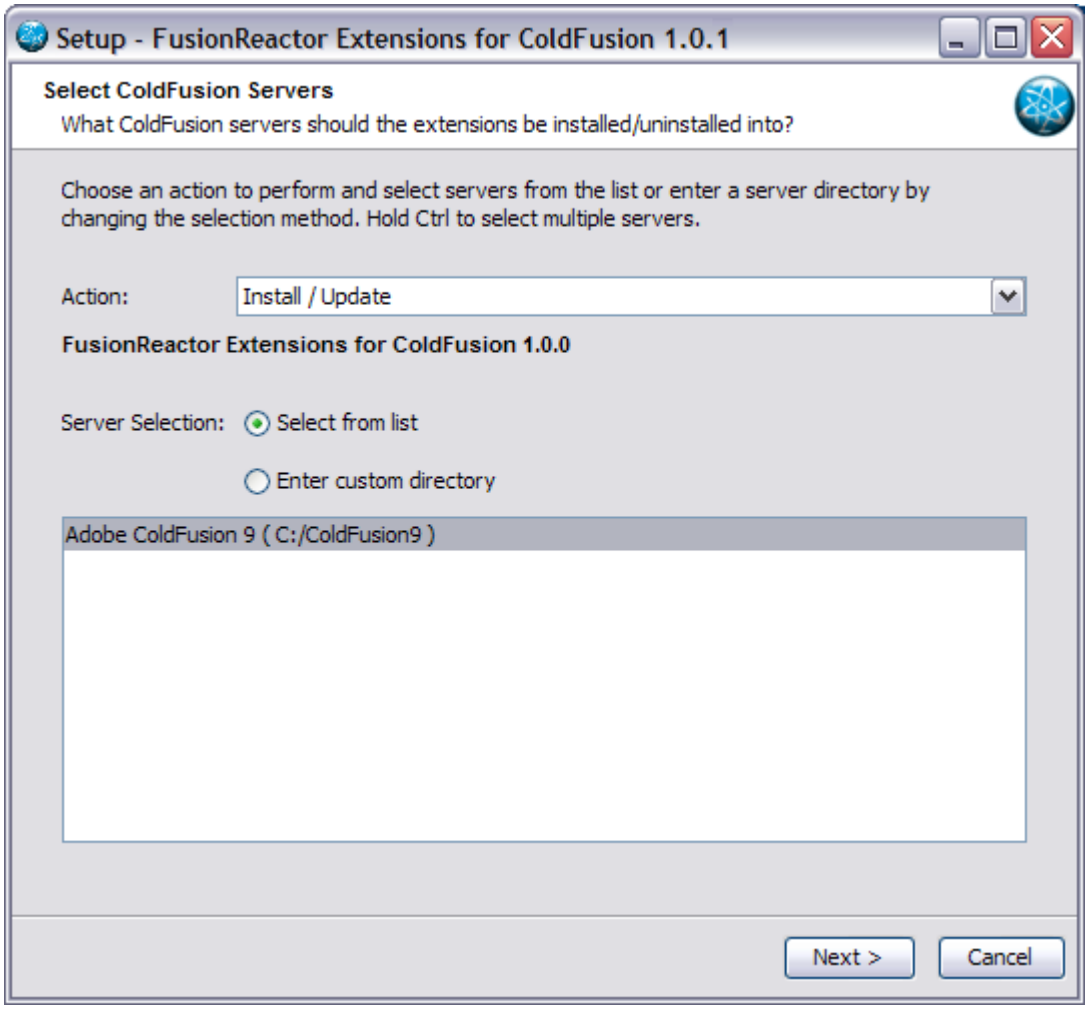


Figure 1: Install / Update FREC into selected servers

The main screen for the Instance Installer (Figure 1) shows a list of the FusionReactor instances which it has found by scanning the system. In order to install the FREC plugin into a FusionReactor instance, select the instance from the list box and make sure that the *'Install / Update'* option is selected from the drop-down list, then click the *'Next >'* button. (It is possible to select multiple instances from the drop down list.)

> **NOTE**
>
> To install FREC on multiple instances of FusionReactor hold down the Ctrl key and click on each one

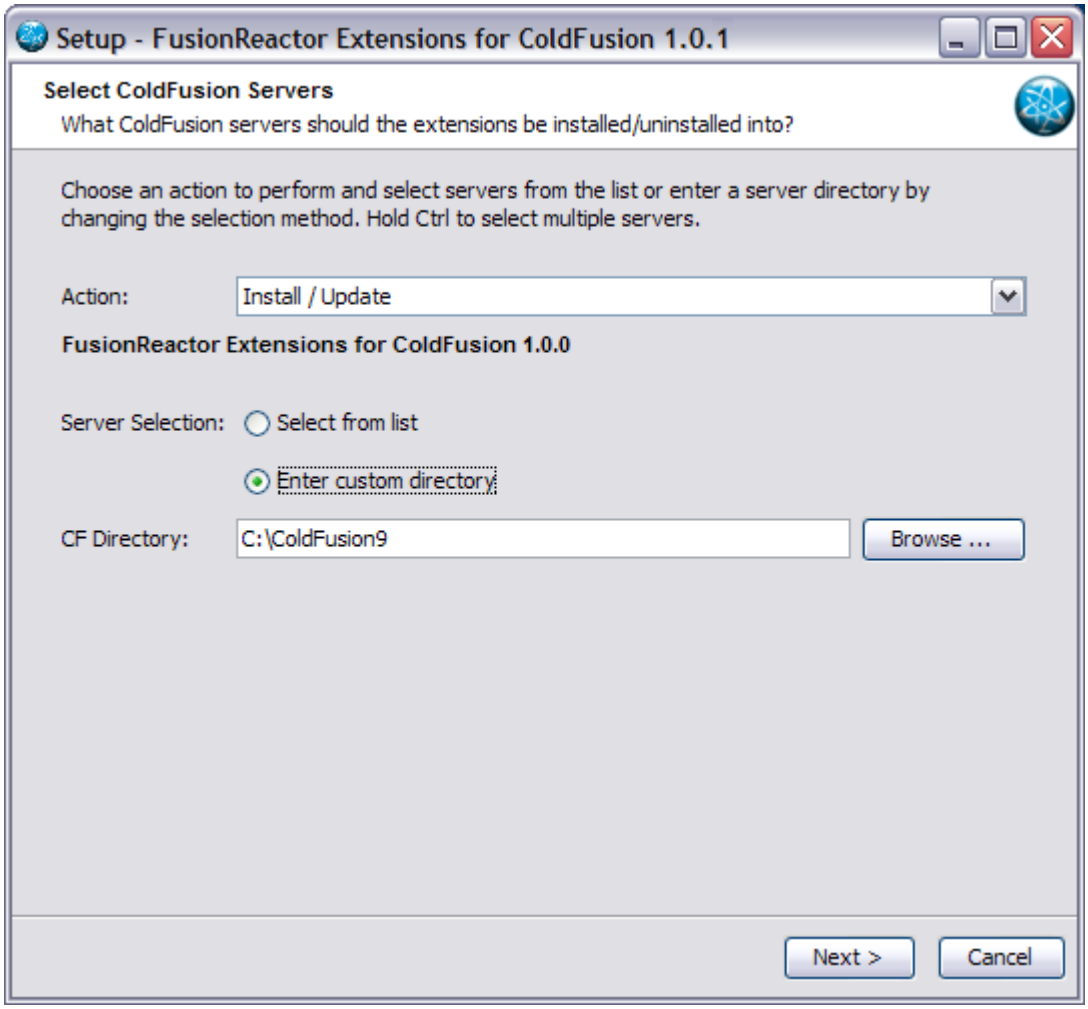
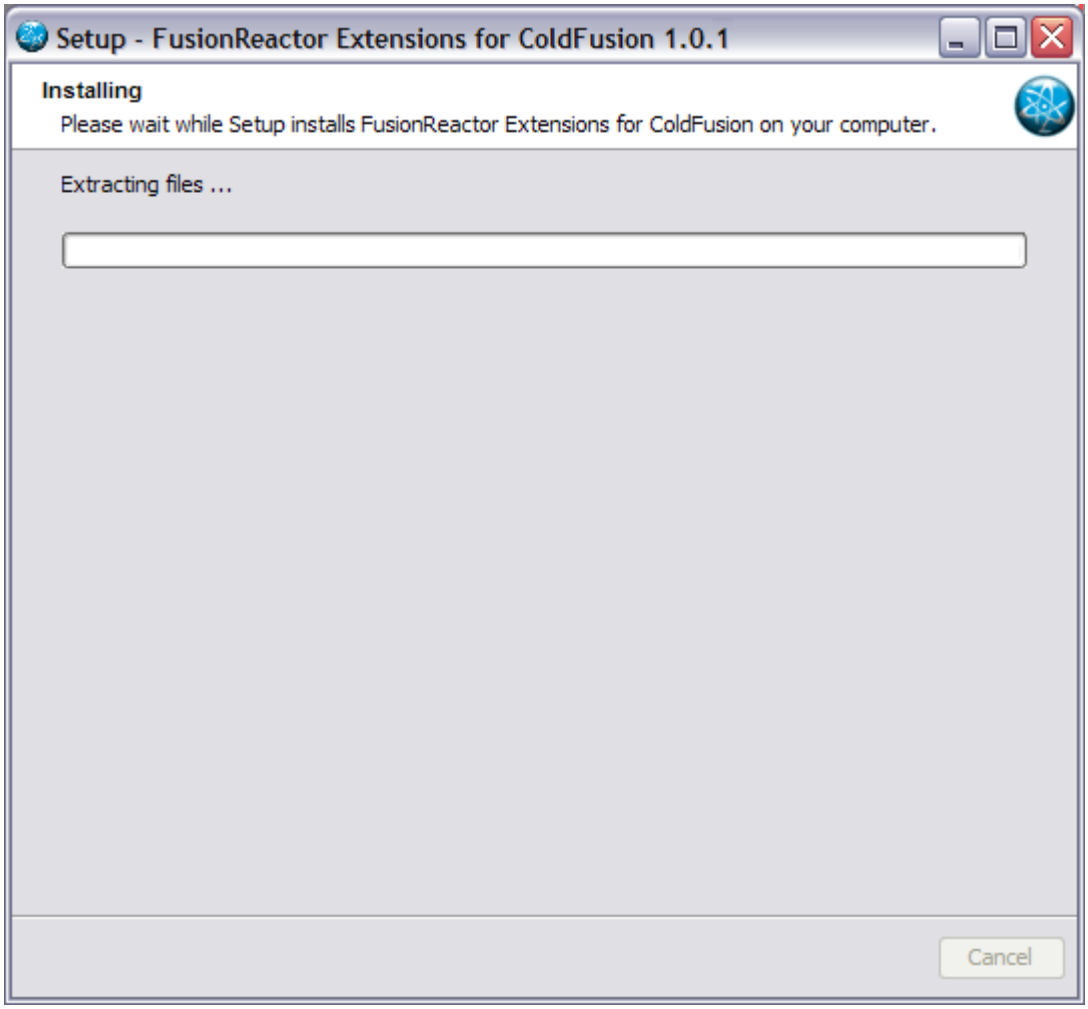


Figure 2: Custom Directory

If you have a FusionReactor instance which isn't listed on this screen then you can select the *'Enter custom directory'* radio button (Figure 2) and enter the directory path in manually. Again, once you have entered the path, click the *'Next >'* button.



Fusion 3: Installing FREC into server instances

The progress screen (Figure 3) shows how long the installation process will take.



Fusion 4: Finish Screen

Once the plugin has been successfully installed into your FusionReactor instance(s) you will see the screen shown in figure 4. Here you will be reminded that your instances will need to be restarted before the plugin will be loaded and become functional. If you wish to install or uninstall other FusionReactor instances then click the *'Go Back to Server Selection'* link. Otherwise click the *'Finish'* button to close the application.

If you want to check the status of the plugin or verify that it installed correctly, navigate to your target FusionReactor instance and select either Plugins -> Active Plugins, or Plugins -> All Bundles from the Table of Contents. You will see a row on both of those screens named *'FusionReactor ColdFusion Plugin'* (Figures 5 and 6 respectively). Remember that you won't see these items listed until you have restarted the target instance after installing the plugin.

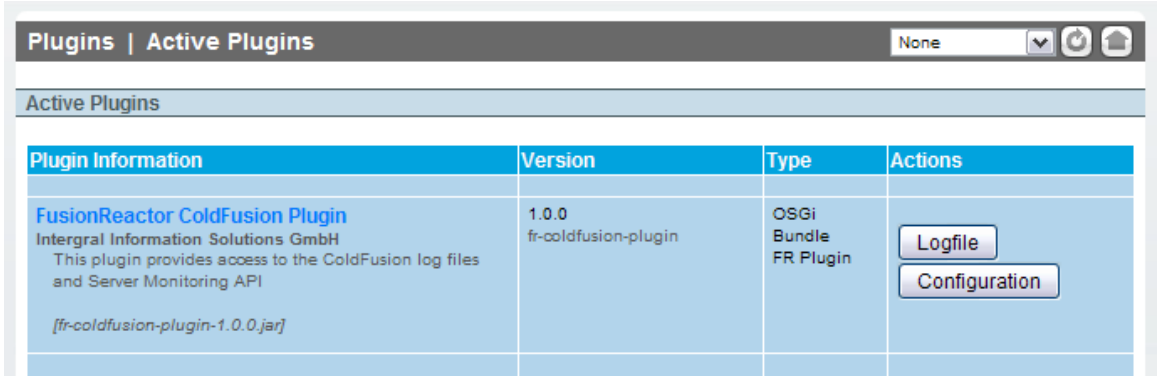


Figure 5: The plugin within active plugins in FusionReactor.

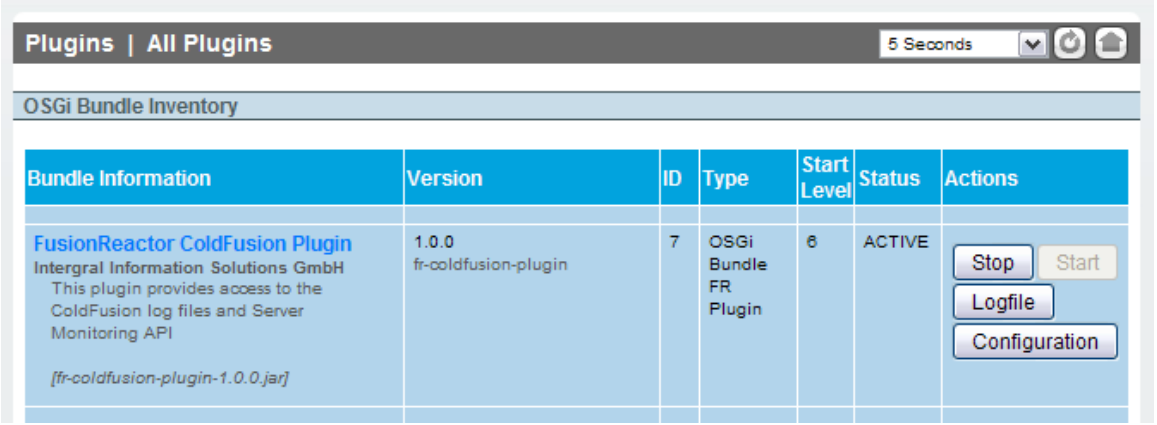


Figure 6: The plugin within all plugins in FusionReactor.

See Also

FREC Configuration Guide

Uninstalling FREC from FusionReactor Instances

In order to remove the FusionReactor Extensions for ColdFusion (FREC) plugin from in instance of FusionReactor you need to run the FREC Instance Installer (eg. FusionReactorColdFusionExtensions -> Manage from the Windows Start Menu). Once you have run it you will see the following page.

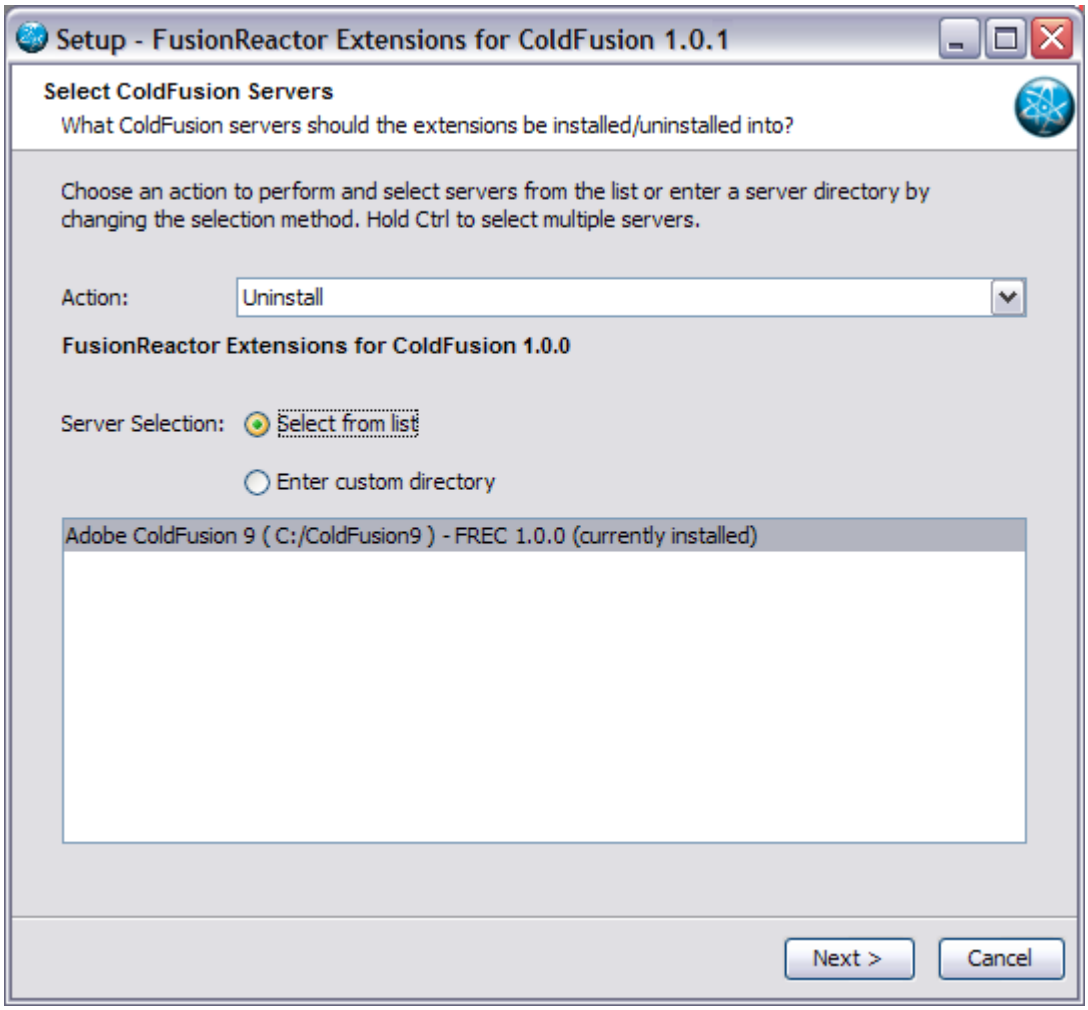


Figure 1: Uninstall List

The main screen for the Instance Installer (Figure 1) shows a list of the FusionReactor instances which it has found by scanning the system. Any instances which have already had the FREC plugin installed into them will include text to indicate this. eg. "FREC 1.0.0 (currently installed)". In order to uninstall the FREC plugin from a FusionReactor instance, select the instance from the list box and make sure that the "Uninstall" option is selected from the drop-down list, then click the "Next >" button.

> **NOTE**
>
> Hold down Ctrl and click on the instances you want to uninstall if you wish to uninstall multiple instances.

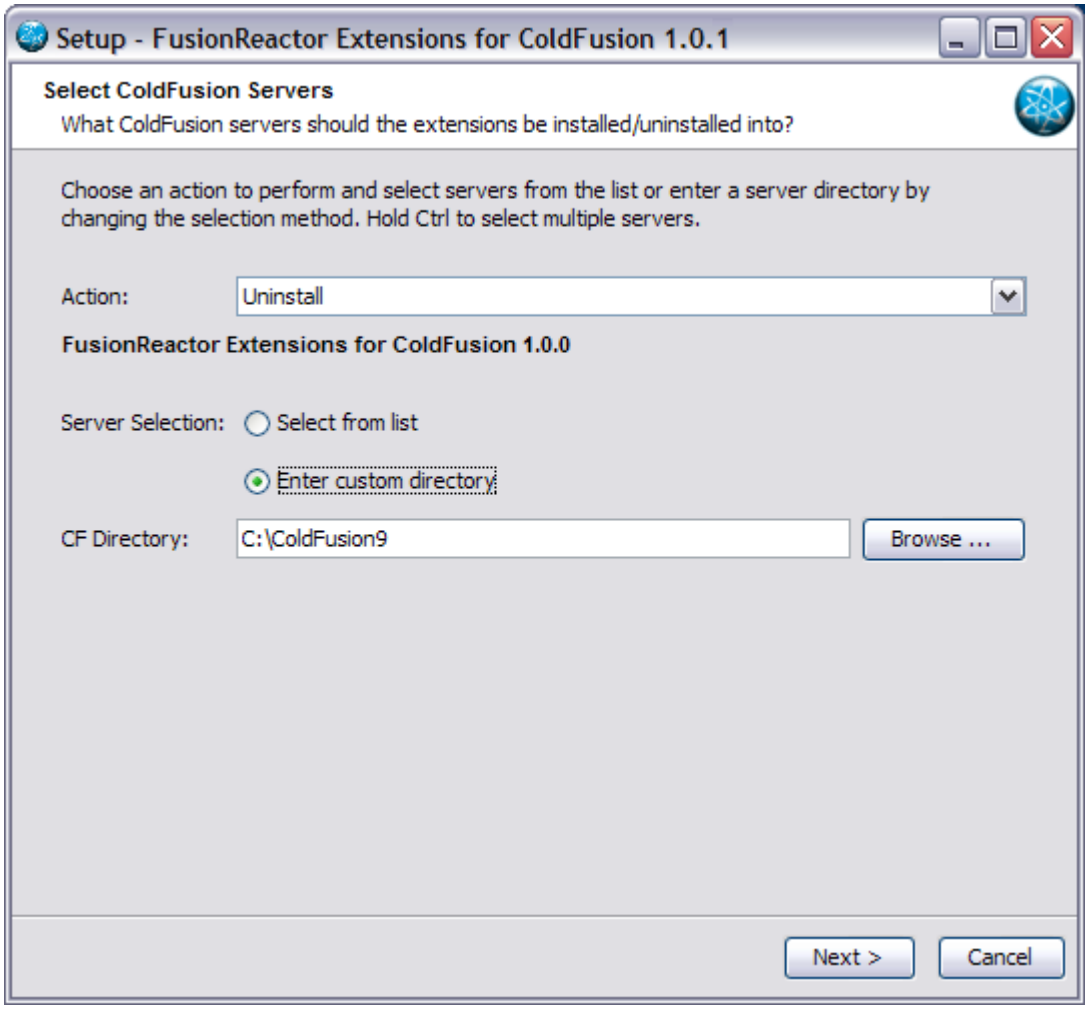


Figure 2: Uninstall Custom Directory

If you have a FusionReactor instance which isn't listed on this screen then you can select the "Enter custom directory" radio button (Figure 2) and enter the directory path in manually. Again, once you have entered the path, click the "Next >" button.

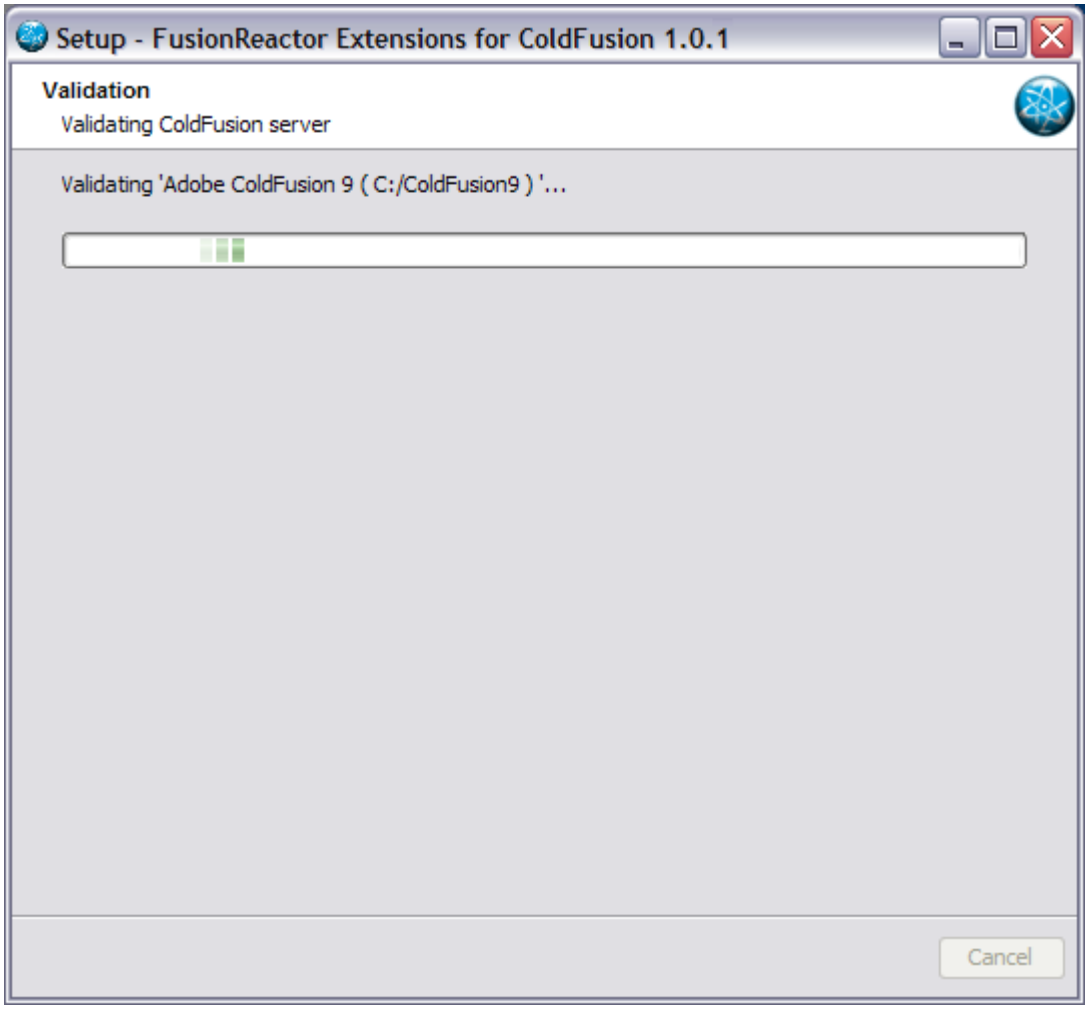


Figure 3: Uninstall Progress

The progress screen (Figure 3) shows how long the uninstall process will take.

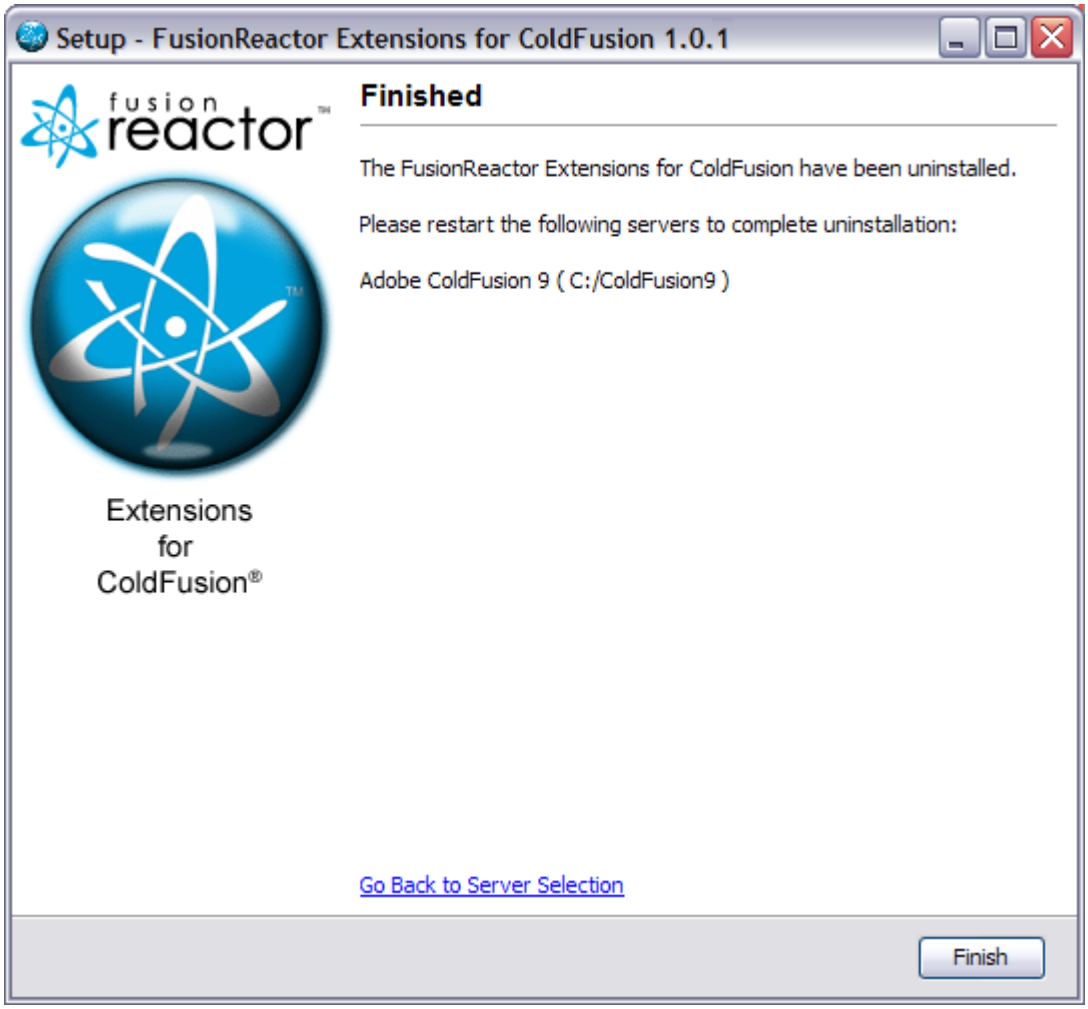


Figure 4: Finished Screen

Once the plugin has been successfully uninstalled from your FusionReactor instance(s) you will see the screen shown in figure 4. Here you will be reminded that your instances will need to be restarted before the plugin actually stops running. If you wish in install or uninstall other FusionReactor instances then click the "Go Back to Server Selection" link. Otherwise click the "Finish" button to close the application.

Uninstalling FusionReactor Extensions for ColdFusion (FREC)

> The FusionReactor Extensions for ColdFusion Instance Installer is completely responsible for the adding or removing of the FREC plugin to/from your FusionReactor instances. If you uninstall the Instance Installer without first uninstalling all plugins then you will no longer be able to administer those plugins.

To uninstall the "FusionReactor Extensions for ColdFusion Instance Installer" select the menu item from the Start Menu or use Add/Remove Programs (Windows platforms.)



Figure 1: Uninstall Welcome Screen

The first screen in the uninstaller is just the Welcome page. Click "Next >" to continue.

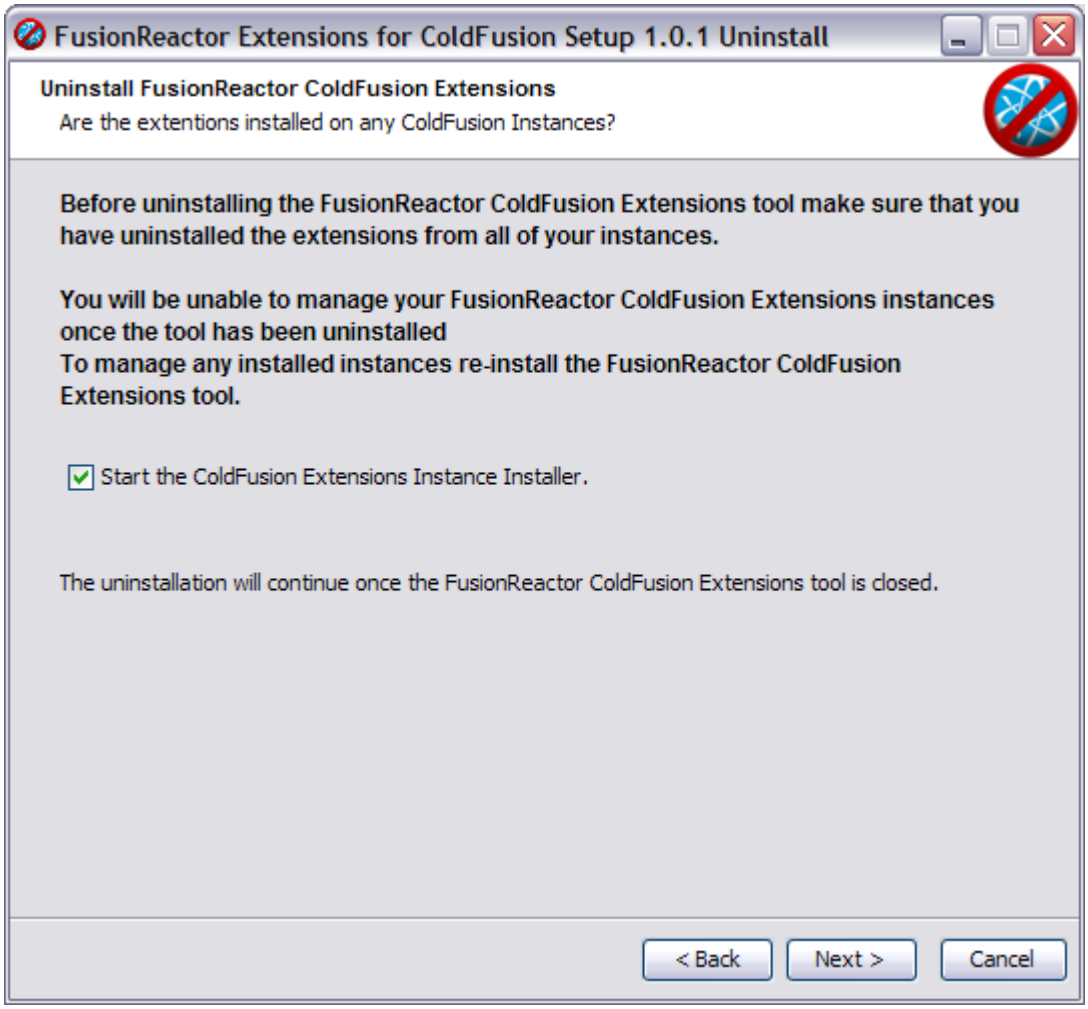


Figure 2: Uninstall Warning

If you have installed the FREC plugin into any instances of FusionReactor then you will see the screen shown in Figure 2. Here you are told that you should make sure all FREC plugins have been removed from FusionReactor instances before continuing. By default, clicking the "Next >" button will open the Instance Installer application where you should remove all of the plugins.

Once you have closed the Instance Installer application, the uninstall process will continue.

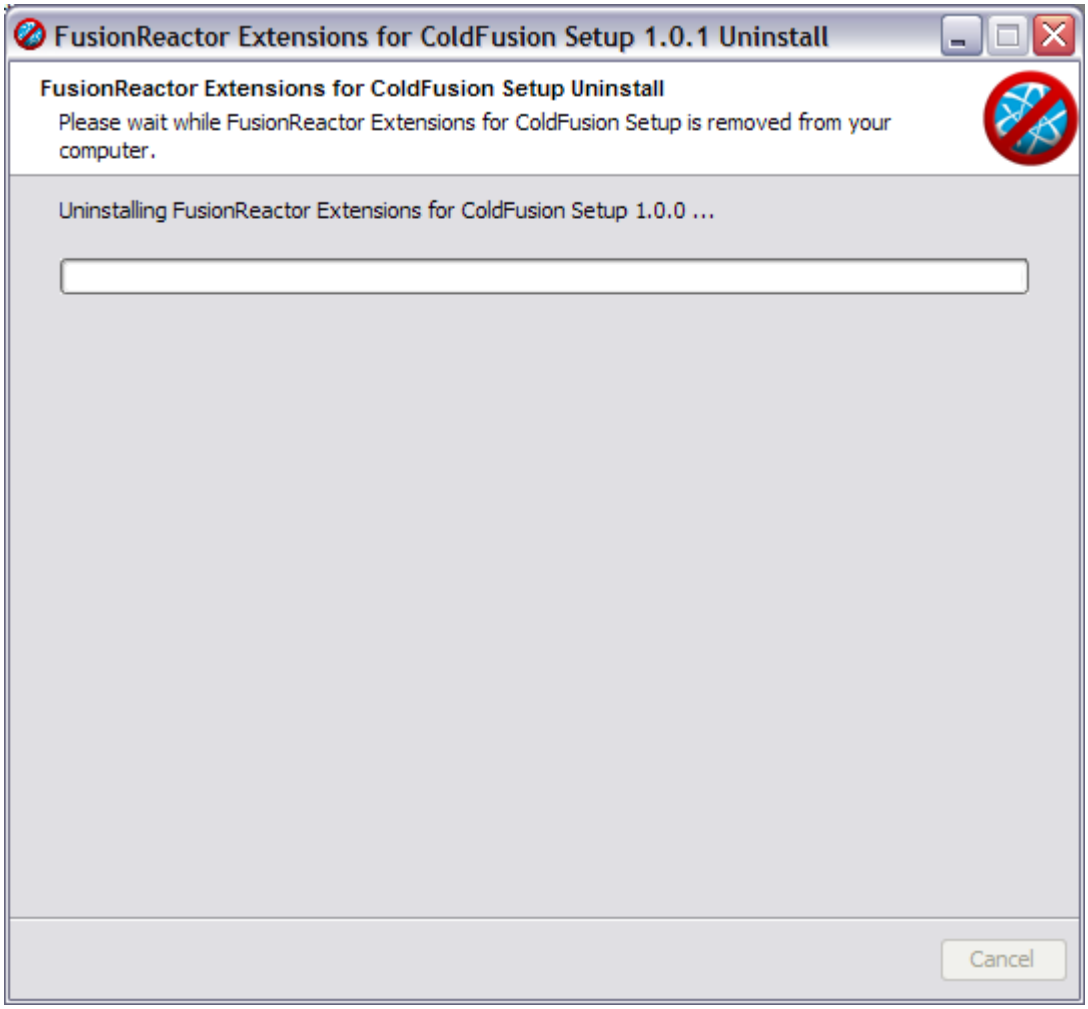


Figure 3: Uninstalling

The third page in the uninstall process is the Progress screen (Figure 3). Wait for the uninstall process to complete.



Figure 4: Finished Screen

Finally you will see the screen shown in Figure 4. Click the "Finish" button to complete the uninstall process and close the window.

FREC Log Files

When FREC is enabled (See How to enable FREC) it will generate extra logs every time FusionReactor performs a Centralized Log Rotation. The following files generated by FREC are as follows:

- dbpoolstats.log
- hitcountstats.log
- realtimestats.log
- plugin-fr-logrotator-plugin.log

Click one of the above links to see a page dedicated to each log file explaining what each of the attributes mean.

CF Database Connection Statistics (dbpoolstats.log)

> **NOTE**
>
> It is important that you enable Monitoring and Profiling in ColdFusion, to do this check out ColdFusion Settings Section in How to enable FREC guide.

This log (dbpoolstats.log) lists the data sources, whether an application on the ColdFusion server is using the data source and the number of connections. The information can be fed through to FusionAnalytics to produce detailed graphical representations of the data. Below is a table explaining each of the attributes in the log. For further information on what FusionAnalytics is and how it can benefit you, click here.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date/Time that the request was made (Milliseconds)	3 [C]	This is the millisecond time at which this request was made.
Get Log Version	4 [D]	This is the version of the log file.
Date/Time that FusionReactor stated (Milliseconds)	5 [E]	The time that the instance of FusionReactor (that is running the plugin) started.
Average Open Connection Count	6 [F]	The average number of open connections that are in the pool since it was last reset.
Average Total Connection Count	7 [G]	The average number of connections that are in the pool since it was last reset.
DNS (Data Source Name)	8 [H]	The name of the data source.
Maximum Number of Connections	9 [I]	The max amount of connections that the pool may hold. This value is -1 if not configured.
Number of open Connections	10 [J]	The amount of connections that are currently open.
Total Connection Count	11 [K]	The amount of connections that are currently in the pool.

Also See

Log Settings
Installing FusionReactor Extensions for ColdFusion (FREC)
Installing FREC into FusionReactor Instances
Uninstalling FREC from FusionReactor Instances
Uninstalling FusionReactor Extensions for ColdFusion (FREC)
FREC Log Files

CF Hit Count Statistics (hitcountstats.log)

NOTE

It is important that you enable Monitoring and Profiling in ColdFusion, to do this check out ColdFusion Settings Section in How to enable FREC guide. Below is a table explaining each of the attributes in the log.

This log (hitcountstats.log) stores hit count related information. Below is a table explaining each of the attributes in the log. The information can be fed through to FusionAnalytics to produce detailed graphical representations of the data. For further information on what FusionAnalytics is and how it can benefit you, click here.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date/Time (Milliseconds)	3 [C]	This is the millisecond time at which this request was started
Get Log Version	4 [D]	This is the version of the log file.
Date/Time that FusionReactor stated (Milliseconds)	5 [E]	The time that the instance of FusionReactor (that is running the plugin) started.
Template Hit Count	6 [F]	The number of hits by a regular HTTP template request.
Delta Template Hit Count	7 [G]	The difference between the current template hit count and the previous template hit count.
Gateway Hit Count	8 [H]	The number of hits by a gateway.
Delta Gateway Hit Count	9 [I]	The difference between the current gateway hit count and the previous gateway hit count.
Flash Hit Count	10 [J]	The number of hits by Flash Remoting.
Delta Flash Hit Count	11 [K]	The difference between the current flash hit count and the previous flash hit count.
RCFC Hit Count	12 [L]	The number of hits from HTTP calls to CFC functions.

Delta RCFC Hit Count	13 [M]	The difference between the current RCFC hit count and the previous RCFC hit count.
WS Hit Count	14 [N]	The number of hits from web services.
Delta WS Hit Count	15 [O]	The difference between the current web services hit count and the previous web services hit count.
Total Hit Count	16 [P]	The total hit count
Delta Total Hit Count	17 [Q]	The difference between the current total hit count and the previous total hit count.
Delta Capture Time (milliseconds)	18 [R]	The difference between the current capture time and the previous capture time.

NOTE

In the fields titled "Delta" the definition describes it as "The difference between". This means the difference between the current data sample and the previous, for example if the current was 10 and the previous was 8 then the delta would be 2 (10 - 8).This sort of information is useful for spotting trends.

Also See

Log Settings
Installing FusionReactor Extensions for ColdFusion (FREC)
Installing FREC into FusionReactor Instances
Uninstalling FREC from FusionReactor Instances
Uninstalling FusionReactor Extensions for ColdFusion (FREC)
FREC Log Files

CF Log Entries (fr-coldfusion-log.log)

The `fr-coldfusion-log` will capture a copy of ColdFusion log entries, for example when making a call to `<cf log>`. This log file will be appened to every time a CF Log event occurs.
The following table describes the captured information.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.

Time (Formatted)	2 [B]	The time on which this log entry was written.
Name of the logger	3 [C]	The name of the logger of the event.
Level	4 [D]	The level, priority, of the logging event.
Message	5 [E]	The message of the logging event.

Example

The file `testFrec.cfm` contains the following:

```
<cflog  file="test" application="no"
text="Test Warning" type="Warning">
```

By running this page a `fr-coldfusion-log.log` is created, containing the following log entry:

```
2011-08-30 16:08:07.421
coldfusion.user.test WARN Test Warning
```

Also See

Log Settings
Installing FusionReactor Extensions for ColdFusion (FREC)
Installing FREC into FusionReactor Instances
Uninstalling FREC from FusionReactor Instances
Uninstalling FusionReactor Extensions for ColdFusion (FREC)
FREC Log Files

CF Realtime and Request Stats (realtimestats.log)

> **NOTE**
>
> It is important that you enable Monitoring and Profiling in ColdFusion, to do this check out ColdFusion Settings Section in How to enable FREC guide.

This log (realtimestats.log) stores real time related information relating to currently active requests. Requests include browser requests, CFC HTTP requests, web services, gateways, and Flash remoting. The information can be fed through to FusionAnalytics to produce detailed graphical representations of the data. For further information on what FusionAnalytics is and how it can benefit you, click here.

The log entries are updated every 5 seconds. Below is a table explaining each of the attributes in the log.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date/Time (Milliseconds)	3 [C]	This is the millisecond time at which this request started.
Get Log Version	4 [D]	This is the version of the log file.
Date/Time that FusionReactor stated (Milliseconds)	5 [E]	The time that the instance of FusionReactor (that is running the plugin) started.
Active Session Count	6 [F]	The current amount of sessions on the server.
Application Scope Size (bytes)	7 [G]	The size of the application.
Average Response Time (milliseconds)	8 [H]	The average server response time.
Cached Query Count	9 [I]	The number of queries that are stored in cache.
CFC Requests Queued	10 [J]	The number of CFC requests (Sent via HTTP) that were queued and are waiting to run.
CFC Requests Running	11 [K]	The number of CFC requests (Sent via HTTP) that are currently running.
CFC Requests Timed Out	12 [L]	The number of CFC requests (Sent via HTTP) that were timed out whilst waiting to run.
CF Thread Queued	13 [M]	The number of CF threads that are queued. NOTE : Value currently set at ZERO
CF Thread Running	14 [N]	The number of CF threads that are running.
Concurrent User Count	15 [O]	The number of users that logged in using the "CFLOGIN" tag.

Flash Remoting Request Queued	16 [P]	The number of Flash Remoting Requests that are queued and waiting to run.
Flash Remoting Request Running	17 [Q]	The number of Flash Remoting Requests that are running.
Flash Remoting Request Time out	18 [R]	The number of Flash Remoting Requests that were timed out whilst waiting to run.
JVM Free Memory (bytes)	19 [S]	The free memory in the Java Virtual Machine.
JVM Used Memory (bytes)	20 [T]	The memory used by the Java Virtual Machine.
Query Cache Hit Ratio	21 [U]	The query cache hit ratio.
Query Cache Size (bytes)	22 [V]	The size of the cache.
Requests Queued	23 [W]	The number of request threads queued (JRun).
Requests Running	24 [X]	The number of requests that are running (JRun).
Requests Timed Out	25 [Y]	The number of request timeouts (JRun).
Request Load	26 [Z]	Request load of the server, displayed as requests per a second.
Server Scope Size (bytes)	27 [AA]	The server scope size.
Session Scope Size (bytes)	28 [AB]	The size of the session scope.
Template Cache Count	29 [AC]	The number of templates stored in the template cache. The template cache is where ColdFusion stores compiled CFM and CFC templates in memory. When a template is executed for the first time, it is compiled to Java bytecode, and then stored in the template cache. As long as the template is unchanged, ColdFusion uses the compiled form of the template stored in the template cache.

Template Cache Hit Ratio	30 [AD]	The hit ratio of the template cache. The cache-hit ratio indicates the number of cache hits in relation to the number of cache misses
Template Cache Size (bytes)	31 [AE]	The total size of the template cache.
Template Requests Queued	32 [AF]	The number of template requests that are queued and waiting to run.
Template Requests Running	33 [AG]	The number of template requests that are running.
Template Requests Timed Out	34 [AH]	The number of template requests that have timed out.
Throttle Queue Size	35 [AI]	The total number of requests queued in the request throttle. Requests are throttled when ColdFusion queues them, because not enough total memory is available to handle them. Requests smaller than the specified limit are not queued or counted as part of the total memory. Requests larger than the specified limit are counted as part of total memory and are queued if the request throttle-memory size of the request is exceeded.
Total Throttle Memory Used (bytes)	36 [AJ]	The total amount of memory that has been used by the throttle.
Web Service Requests Queued	37 [AK]	The number of Web Services that are queued and waiting to run.
Web Service Requests Running	38 [AL]	The number of Web Services requests that are currently running.
Web Service Requests Timed Out	39 [AM]	The number of Web Services that timed out whilst waiting to run.
Memory Monitor State	40 [AN]	If Memory Monitoring is enabled in ColdFusion then this value will be 1, if not it will be 0. To enable this check out How to enable FREC.

Profiling State	41 [AO]	If Profiling is enabled in ColdFusion then this value will be 1, if not it will be 0. To enable this check out How to enable FREC.

More detailed information relating to these metrics can also be found within the Adobe ColdFusion documentation.

Also See

Log Settings
Installing FusionReactor Extensions for ColdFusion (FREC)
Installing FREC into FusionReactor Instances
Uninstalling FREC from FusionReactor Instances
Uninstalling FusionReactor Extensions for ColdFusion (FREC)
FREC Log Files

Configure FusionReactor ColdFusion Plugin (FREC)

The FusionReactor ColdFusion Plugin (FREC) plugin can be configured for each individual instance it has been installed into. Navigate to the **"Active Plugins"**, or **"All Plugins"**, pages available in the **"Plugins"** menu item in the FusionReactor interface of the FusionReactor instance FREC has been installed into.

In the list of plugins, locate the plugin named **"FusionReactor ColdFusion Plugin"**, see Figure 1.



Figure 1: FusionReactor ColdFusion Plugin in Active Plugins

Figure 1: To Configure the FREC plugin, click the button **'Configuration'**. A new page will be displayed, Figure 2.

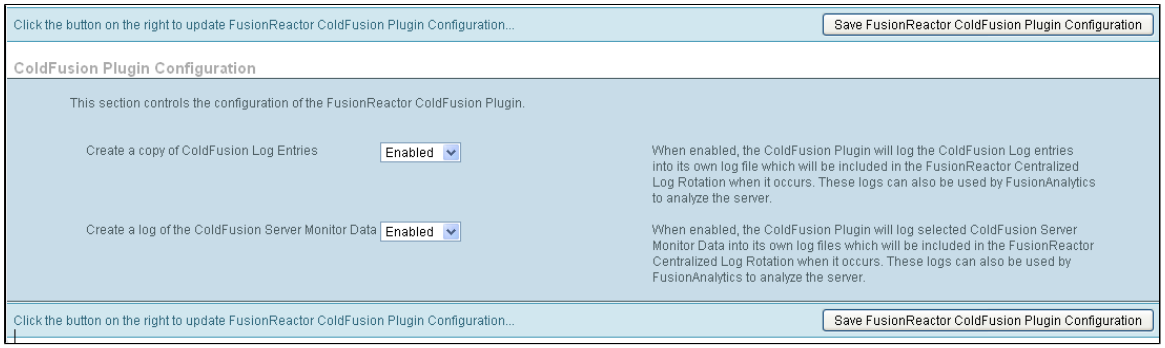


Figure 2: Configure FusionReactor ColdFusion Plugin

Figure 2: The following configurations options are available to Enable or Disable.

Configuration	Description
Create a copy of ColdFusion Log Entries (**Available in ColdFusion 7, 8, 9 and 10**)	When enabled, FREC will save a copy of the ColdFusion log entries (such as calls to `<cflog>`) to the following log file located in the FusionReactor instance's log directory. • fr-coldfusion-log.log

Create a log of the ColdFusion Server Monitor Data (**Available only in ColdFusion 8,9 and 10**)	When enabled, FREC will save selected ColdFusion Server Monitor Data to the following log files located in the FusionReactor instance's log directory. • dbpoolstats.log • hitcountstats.log • realtimestats.log

These log files will be included in the Centralized Log Rotation, which archives and pushes log files to FusionAnalytics.

> **NOTE**
>
> Enabling these options will allow FusionAnalytics (if configured in the FusionReactor instance) to analyze these logs, giving you a better insight into how your server is performing. Disabling them is not recommended.

Click the **'Save FusionReactor ColdFusion Plugin Configuration'** button to save the configuration, changes will be applied immediately.

Also See

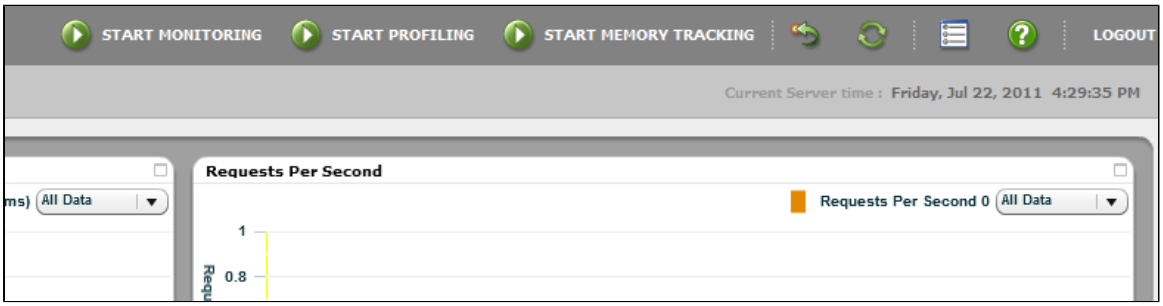
FREC Log Files

How to enable FREC

Once FREC is installed, you may have to turn on specific monitoring within the ColdFusion Server Monitor in order to see reasonable data in the generated log files. This page details how to enable the different monitoring levels with the **ColdFusion Server Monitor**. If a required level is not enabled, FREC will zero that log field, to ensure each log line remains integral.

ColdFusion Settings

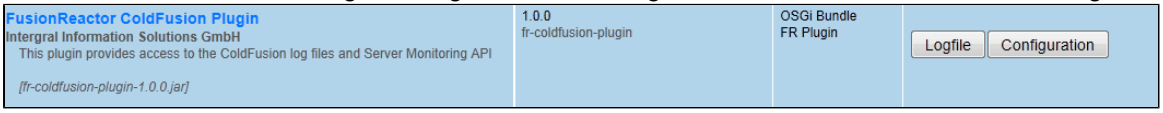
In the ColdFusion contents page, click Server Monitoring > Server Monitor. You will be presented with a page with some graphs and 3 buttons at the top, their names being "START MONITORING", "START PROFILING" and "START MEMORY TRACKING". Below is a screen shot that shows this.



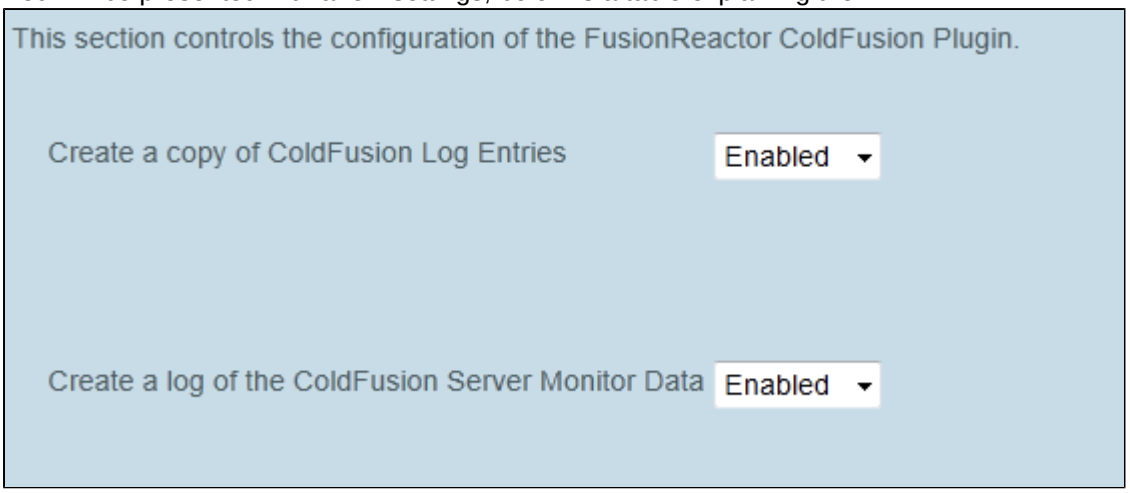
Using the reference pages for each FREC-produced log file, select one or more levels to turn on by clicking the button. Once enabled, the button will turn red.

FusionReactor Settings

In a FusionReactor instance go to Plugins > Active Plugins > FusionReactor ColdFusion Plugin > Configuration.



You will be presented with a few settings, below is a table explaining them.



Configuration	Description
Create a copy of ColdFusion Log Entries (**Available in ColdFusion 7, 8, 9 and 10**)	When enabled this the plugin will log ColdFusion Log entries and include them in the FusionReactor Centralized Log Rotation. These logs can also be used by FusionAnalytics to analyze the server. Enabling this will create the following logs. • plugin-fr-logrotator.log
Create a log of the ColdFusion Server Monitor Data(**Available only in ColdFusion 8, 9 and 10**)	When enabled the plugin will log selected ColdFusion Server Monitor Data into its own log files, these log files will be included in the FusionReactor Centralized Log Rotation when it occurs. They can also be used by FusionAnalytics to analyze the server. Enabling this will create the following logs. • dbpoolstats.log • hitcountstats.log • realtimestats.log

Enable for of these for the plugin to work.

Summary

After enabling the settings mentioned above in ColdFusion and FusionReactor the logs will be created every time FusionReactor performs a Centralized Log Rotation.

Also See

Log Settings
Installing FusionReactor Extensions for ColdFusion (FREC)
Installing FREC into FusionReactor Instances
Uninstalling FREC from FusionReactor Instances
Uninstalling FusionReactor Extensions for ColdFusion (FREC)
FREC Log Files

FusionReactor Heartbeat Plugin

Description

Constantly logs whether or not a specific ColdFusion instance is up (running).

> **NOTE**
>
> Increasing the value of the Heartbeat Metrics Fetch Interval will lower the refresh rate of the metric data shown on the Enterprise Dashboard. Decreasing this value may cause heartbeat failures. It is NOT recommended to set this value below 5 seconds.

Configuration

This plugin can be configured in Enterprise Settings.

Configuration	Description
Heartbeat Metrics Fetch Interval (seconds)	This interval tells FusionReactor how long to wait between each fetch of a monitored instance's metric data (its heartbeat). The heartbeat contains the data that is shown on the Enterprise Dashboard.
Heartbeat Timeout (seconds)	This timeout governs how long FusionReactor will wait for a heartbeat from a monitored server instance. If no heartbeat is received within this time or a network connection failure occurs, FusionReactor will count it as a heartbeat failure.
Heartbeat Failure Threshold	This setting tells FusionReactor how many times the heartbeat may fail on a monitored server instance before considering it to be DOWN. FusionReactor will then report the server instance as [CONN] on the Enterprise Dashboard and will run the self-healing script for the instance (if configured). If you are receiving monitoring emails and log messages for instances known to be available, try increasing this value. A value of zero will cause the first missing heartbeat to signal the server instance is DOWN.

Log File - heartbeat.log

The plugin generates a log file called **heartbeat.log**, this contains heartbeat information that could be very useful when diagnosing server problems.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date the log was written.
Time (Formatted)	2 [B]	The time the log was written.
Date/Time (ms)	3 [C]	The number of seconds since epoch.
Version	4 [D]	Version number of this log entry.

Server Start up Time (ms)	5 [E]	Start up time of the application server in milliseconds (ms) since epoch.
Beat	6 [F]	If the server is up the value is 1, if not it will be 0. When your server starts up it read read 1, any other entry after that will read 0.
Up	7 [G]	If the server has started up this value will be 1, if not it it will be 0.
Down	8 [H]	If the server is shutdown this value will be 1, if not it will be 0.
Dirty Shutdown / Start up	9 [I]	If a dirty shutdown or start up occurred in a server this value will be 1 when the server next starts.
Date/Time	9 [J]	The date and time the entry was written according to ISO 8601 standards.

Status Log - plugin-fr-heartbeat-plugin.log

The Heartbeat plugin creates a log file called plugin-fr-heartbeat-plugin.log which contains information on the status of the plugin, when it started, stopped, if there were any problems. which contains information on the status of the plugin, when it started, stopped, if there were any problems. Below is a table explaining what all the attributes in the log file mean.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information, where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.

Also See

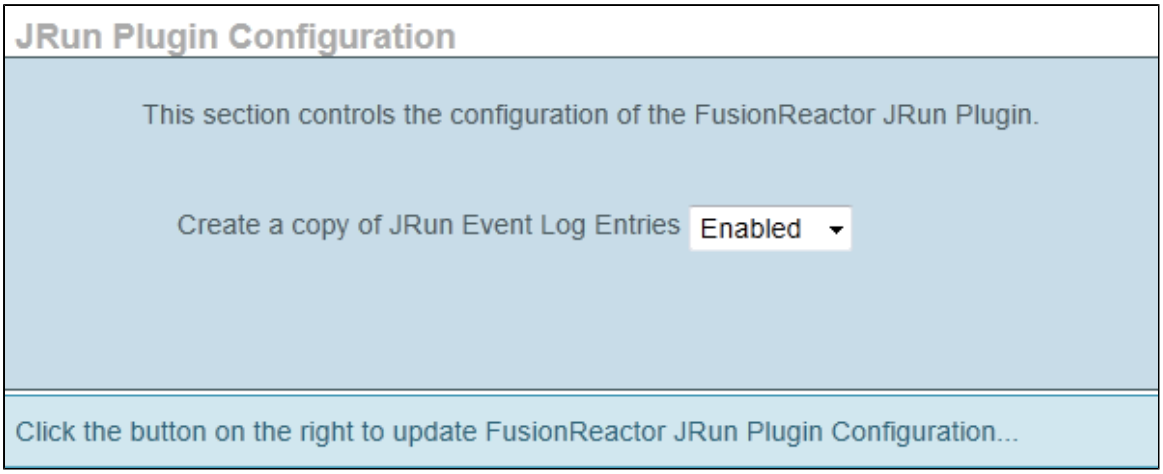
FusionReactor JRun Plugin

Description

This plugin allows you to access the JRun log files.

Configuration

This page allows you to enable or disabled the creation of JRun Event Log Entries. Enabling this will allow the plugin to log JRun server events which are then included in the FusionReactor Centralized Log Rotation when it occurs. The logs can also be used by FusionAnalytics to analyze the server. The option you choose does not alter the JRun Log Directory setting.



Log File - fr.jrun-event.log

This log is called fr.jrun-event.log. **Don't** confuse this with that status log, the status log displays the state of the plugin (on, off and error) where as this describes what the plugin is doing (starting applications etc).

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date (Short)	3 [C]	Displays a short version of the date (DD/MM).
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.
Time (Short)	5 [E]	The time the entry was written.

Entry	6 [F]	The entry that was written. for example the entry might state that a servlet was started up.

Status Log - plugin-fr-jrun-plugin.log

This plugin creates a log file called **plugin-fr-jrun-plugin.log** which contains information on the status of the plugin, when it started, stopped, if there were any problems. Below is a table explaining what all the attributes in the log file mean.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information, where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.

Also See

Log Settings
FusionReactor Plugins

FusionReactor Log Rotator Plugin

Description

Manages the maintenance of log files within FusionReactor. Features include:-

- Two types of mode - *Centralized Archive and Rotation* and *Per-Log Rotation* mode
- Rotates log files on a given schedule (either at a given time every day, or periodically)
- Transfer of rotated log collection as a zip file to a central location
- Management of log files (purge) based on age or size target (oldest first)
- Integration with FusionAnalytics Connector

Configuration

This configuration page allows you to select if logs are archived and transferred to FusionAnalytics or just archived and kept in the archive directory in FusionReactor.

Logs that are just archived can be located at /FusionReactor/instance/[your_instance_name]/log.

The following data can be archived and transferred or just archived:

- CPU/Thread Log
- Memory Detailed (Fine)
- Memory Summary (Coarse)
- Reactor Log
- Request Log
- Resource Log
- Plugin Log
- Crash Protection Log
- JDBC Log
- System Heartbeat
- Default

Log File

No log files about captured data are generated from this plugin.

Status Log - plugin-fr-logrotator.log

This plugin creates a log file called **plugin-fr-logrotator.log** which contains information on the status of the plugin, when it started, stopped, if there were any problems. Below is a table explaining what all the attributes in the log file mean.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information, where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.

Also See

Log Settings
FusionReactor Plugins

FusionReactor Notification Generator Plugin

Description

The Notification Generator plugin allows you to create notifications whilst logged into FusionReactor. The notification is only available in the instance in which it was created in. For example if you created a notification inside an instance called cfusion.cfmx10 then the notification would only be available inside that instance and no other instances i.e. default.tomcat. See the Notifications page for information on what a notification is.

Usage

To create a notification navigate to FusionReactor > Create Notification. You will be presented with three fields, these fields are described below.

Field	Description
Notification Title	The notification title, for example it might be "The server is being shutdown"
Notification Type	The type of notification. Types include Information, Warning or Error. This allows you to communicate the intent of the notification to observers. For example if you were restarting your server you would set the Notification Type to Information.
Notification Message	The notification message, for example it might be "The server is being shutdown for maintenance"

Once you are satisfied click the button titled Create Notification and the notification will be created.

Scenario

I want to shutdown a server to apply a patch and quickly inform others of this. To do this I need to create a notification, that notification will be pushed to users who are observing FusionReactor via the web browser, AIR Enterprise Dashboard, iOS App or the Android App.

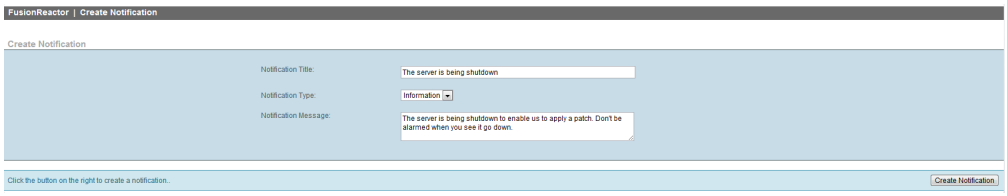


Figure 1: Notification Generator page

Figure 1: The notification details have been input.



Figure 2: Notifications page

Figure 2: The notification has been created. Users observing FusionReactor via the web browser, AIR Enterprise Dashboard, iOS App or the Android App will receive the notification.

Status Log - plugin-fr-notification-generator-plugin.log

The Notification Generator plugin creates a log file called plugin-fr-notification-generator-plugin.log which contains information on the status of the plugin, when it started, stopped and if there were any problems. Below is a table explaining what all the attributes in the log file mean.

Field Name	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Level	3 [C]	Log Level basically means the nature of the log, for example INFO means general information, where as ERROR means an error is being logged. SEVERE means something really bad has happened, this will be typically followed by the a generated exception.
Log message	4 [D]	Contains a message of an event that occurred such as logging that has started.

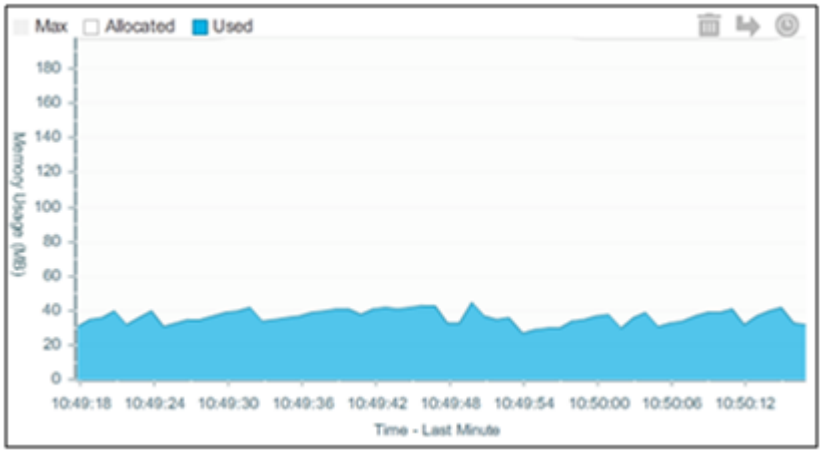
Also See

Notifications

Memory Protection

As an example, we set up a load test to simulate a fairly constant load on the system.

You can see in the System Metrics page that the memory displayed in the **Memory Usage** graph is reasonably constant. If we observe the system for an hour, we would get an even better picture of how the memory is used. For even longer periods, we can set the **Resource Settings -> Resource Sampling** interval to 60000ms (1 minute), with a long history size. We can then view this data in the **Resources -> Memory Graph** page. Using a longer time span allows us to get a picture of time-dependent fluctuations such as peak load and scheduled tasks.



Using the data on the Memory Usage graph, we can see that the steady state for our test is around 40MB. Allowing a 100MB overhead for scheduled tasks and peak-time usage, this would give an acceptable operational peak of 140MB, with a remaining “emergency” margin of 60MB.

We might then set Crash Protection **Memory Protection** to **30%** (60MB margin = 30% of the total memory).

The system should then be left in **Notify** mode for a period of time. While running in this mode, adjustments can be made to the protection value based on the notification email. When you are satisfied that the value you have represents the upper margin for the system, change the notify mode to **Abort and Notify** or **Queue and Notify**.

Methodology in OpenOffice

1. Ensure **Requests -> Request Settings -> Request Logging** is **Enabled.**
2. Allow the system to run in its steady state for a while (the longer the better).
3. Load the **request log** (/FusionReactor/instances/NAME/log/request-0.log) into OpenOffice Calc
    a. Open the file as a **Text CSV** file, with options: **Separated by: Space**, **Text delimiter: “**
4. Create a new sheet (**Insert -> Sheet**), naming the new sheet **Sheet2.**
5. Switch back to **Sheet1**, where the data is, and select **Edit -> Select All**.
6. Select **Data -> Filter -> Standard Filter**, setting **Column F** equal to **COMPLETED**. This will remove any other request statuses from the list. Click the **More** button*,* and select **Copy Results To**. Enter **$Sheet 2.$A$1** in the destination box and click **OK.**
7. When the copy is complete, you should have only the **COMPLETED** entries in Sheet 2.
8. In a free cell which is not in column , enter the formula: **=AVERAGE(L1:L65535)+STDEV(L1:L65535)**
    a. This value is the upper limit at **one standard deviation away from the average.**
9. Statistically,this may still produce too many false positives for Crash Protection, so we may want to use 2, 3 or even 4 standard deviations:
    a. **=AVERAGE(L1:L65535)+4*STDEV(L1:L65535)**

As with the gut feeling approach, this number can only ever be a starting point for tuning: you are bound to get some exceptions, and these should help you decide how to adjust this value.

Restriction

In the detailed example above, we've allowed a margin in the operational memory estimation to account for periodic rises in website usage, and scheduled tasks. Scheduled tasks often run at periods of low demand, and are typically resource-intensive. Search indexing, reporting and cleanup operations often run in this batch mode, and often consume memory and CPU cycles. These tasks also tend to run for long periods of time.

Once you have a picture of these tasks, you can add them as **CP Restrictions** (Crash Protection -> CP Restrictions). Ensure that **CP Settings -> CP Restrictions -> Restrictions** is **Enabled**, and the mode is **Ignore Matching Requests**. CP will then ignore any requested pages set in the **CP Restrictions** page.

Once you have added your scheduled tasks as restrictions, the operational values for **Timeout, Memory** and **Re quest Protection** can be adjusted, since they will no longer apply to the exceptional cases.

Timeout Protection and Request Protection

The **System Metrics** charts may also be used to generate values for **Timeout Protection** (using the **Average Request Time** graph), and **Request Protection** (using the **Request Activity** graph).